

Galaxy® Aurora LS Series RAID Storage System Configuration and System Integration Guide



Version 2.1 February 2011



Rorke Data Inc
7626 Golden Triangle Drive
Eden Prairie MN 55344-3732

952 829 0300

Sales@rorke.com

techsupport@rorke.com

This manual is preliminary and under construction and only applies to the Galaxy® Aurora LS product. Contact Rorke Tech support for specific technical information regarding this manual.

Version 1.0	March 20, 2009
Version 1.1	July 22, 2009
Version 2.0	December 10, 2009
Version 2.1	February 17, 2011

Table of Contents

Copyright 2009	6
Disclaimer	6
Trademarks	6
Notices	6
SAFETY PRECAUTIONS	7
CONVENTIONS	8
Galaxy® Aurora LS EOS / NumaRAID Updates	9
1.0 Introduction and Overview	11
1.1 Product Specifications	11
1.1.1 Overview	11
1.1.2 Basic Features and Advantages	11
1.2 Model Variations	13
1.2.1 Galaxy® Aurora LS Model Descriptions	13
1.3 Product Description	14
1.3.1 Description of Physical Components	14
1.3.2 Component specifications	16
1.3.3 RAID storage specifications	17
1.3.4 Embedded OS features	17
1.4 Mounting / Securing Aurora LS	17
1.4.1 Rack Mounting the Aurora LS	17
1.4.2 Installation Sequence	18
1.4.2.1 Ball Bearing Slide Rail Rack Installation	19
2.0 Basic Setup	23
2.1 Drive integration and Cable Connections	23

2.1.1	Indicators and switch descriptions Figure 2.1	23
2.1.2	Installing drives into the Aurora LS Figure 2.2.....	24
2.1.3	Connecting Cables Figure 2.3.....	25
2.2	Configuration Setup.....	26
2.2.1	Setting up Ethernet Connectivity on a Windows Client	26
2.2.2	Installing Fibre Channel HBA and drivers on Aurora LS Clients	27
2.2.3	Linux Client RAID Connections and LUN Preparation.....	28
2.2.4	Windows Client RAID Connections and LUN Preparation	31
2.2.5	Apple OSX Client RAID Connections and LUN Preparation.....	38
2.3	Remote Administration	43
2.3.1	Using a Browser and Logging into the Aurora LS.....	43
3.0	Aurora LS GUI Detailed Operations	44
3.1.0	GUI Menu Details and Functions	44
3.1.1	Main GUI screen page details and Quick Start functions	44
3.1.2	RAID Creation, Status, and other RAID configuration information	47
3.1.3	RAID Details	49
3.1.4	Scan / Performance Results	51
3.1.5	LUN Details.....	53
3.1.6	CONFIG Details	54
3.1.7	TRACE Details	56
3.1.8	USER Details.....	59
3.1.9	PARAM Details	61
3.1.10	DATARATE Details.....	65
3.1.11	SLOT Details	69
3.1.12	SENSOR Details	73
3.1.13	ADAPTER Details.....	74
4.0	Troubleshooting Aurora LS	75

4.1	Chassis Status Indicators	75
4.2	GUI status indicators	76
4.3	Power System.....	76
4.4	Using GUI for FAN problems	77
4.5	Using GUI for Power Supply problems	77
4.6	DC Power Distribution problems.....	77
4.7	Chassis Problems.....	77
4.8	Motherboard problems.....	79
4.9	Drive Backplane problems	81
4.10	Boot device problems	81
4.11	Data Drive problems	81
4.12	SAS HBA problems	81
4.13	SAS Host connectivity issues	82
4.14	Fibre HBA problems.....	82
4.15	Fibre Host connectivity issues.....	83
4.16	Troubleshooting Aurora LS's Client Related Problems	83
	Fibre Based Clients.....	83
4.17	Using IPMI to diagnose problems	85
5.0	Application / Technical / Customer Notes.....	89
5.1	Windows Infiniband Performance Tuning.....	89
5.2	Additional Administration Functions	98
	System Information.....	98
	IP Address Firewall	99
	Default to NumaRAID GUI after Login	100
	Make the NumaRAID GUI a Little Faster	100
	Find the IP Addresses of Other NumaRAID(s) on the Network	101
	Adding/Deleting/Changing Webmin Users	101

Changing Passwords	101
Run a CLI command from Webmin	102
Change the Network Host Name	102
See and Control SMART for the Boot Device.....	102
Setting System Time or Timezone	103
Logging Out.....	103
5.3 Fibre Channel Switch Zoning	105
5.4 Infiniband Switch Configurations.....	106

Copyright 2009

This Edition First Published 2009 All rights reserved. This publication may not be reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any language or computer language, in any form or by any means, electronic, mechanical, magnetic, optical, chemical, manual or otherwise, without the prior written consent.

Disclaimer

Rorke Data makes no representations or warranties with respect to the contents hereof and specifically disclaims any implied warranties of merchantability or fitness for any particular purpose. Furthermore, Rorke Data reserves the right to revise this publication and to make changes from time to time in the content hereof without obligation to notify any person of such revisions or changes. Product specifications are also subject to change without prior notice.

Trademarks

Rorke Data and the Rorke Data logo are registered trademarks of Rorke Data, Inc. Rorke Data and other names prefixed with “Aurora LS” and “Galaxy” are trademarks of Rorke Data, Inc. in the United States, other countries, or both.

Infiniband is a registered trademark of System I/O, Inc.

LSI and SAS-1068e are registered trademarks of LSI Logic, Inc.

Mellanox, ConnectX, and Infinihost are registered trademarks of Mellanox, Inc.

Microsoft, Windows, Windows XP, Windows 2003, and Windows Vista are registered trademarks of Microsoft Corp.

OFED is a registered trademark of the Open Fabrics Alliance.

All other names, brands, products or services are trademarks or registered trademarks of their respective owners.

Notices

The content of this manual is subject to change without notice. Although steps have been taken to create a manual which is as accurate as possible, it is possible this document may contain inaccuracies or that changes have been made to the system.

Safety Precautions

Precautions and Instructions

- The Aurora LS weighs over 80 pounds. It is recommended that 2 people are required to properly move and mount it due to its size and weight.
- Be sure that the rack cabinet into which the subsystem chassis will be installed provides: sufficient strength and stability and ventilation channels and airflow circulation around the subsystem.
- **INSTALL AURORA IN RACK MOUNTING BEFORE INSTALLING DISK DRIVES**
- The Aurora LS RAID subsystem will come with up to twelve (12) drive bays. NOTE: a 13th drive slot is available but is not used in the present configuration. This drive slot, the furthest right slot for all drives, must have an empty drive canister in its place. Leaving any of these drive bays empty will greatly affect the efficiency of the airflow within the enclosure, and will consequently lead to the system overheating, which can cause irreparable damage.
- Prior to powering on the subsystem, ensure that the correct power range is being used.
- If a disk or power supply module fails, leave it in place until you have a replacement unit and you are ready to replace it.
- Airflow Consideration: The subsystem requires an airflow clearance, especially at the front and rear.
- Handle subsystem modules using the retention screws, extraction levers, and the metal frames/faceplates. Avoid touching PCB boards and connector pins.
- To comply with safety, emission, or thermal requirements, none of the covers or replaceable modules should be removed. Make sure that during operation, all enclosure modules and covers are securely in place.
- Provide a soft, clean surface to place your subsystem on before working on it. Servicing on a rough surface may damage the exterior of the chassis.
- If it is necessary to transport the subsystem, repackage all disk drives separately. If using the original package material, other replaceable modules can stay within the enclosure.
-

ESD Precautions

Observe all conventional anti-ESD methods while handling system modules. The use of a grounded wrist strap and an anti-static work pad is recommended. Avoid dust and debris or other static-accumulative materials in your work area.

Conventions

Naming

From this point on and throughout the rest of this manual, the Aurora LS series is referred to as simply the “Aurora LS”, “subsystem” or the “system.”



Important Messages

Important messages appear where mishandling of components is possible or when work orders can be mis-conceived. These messages also provide important information associated with other aspects of system operation. The word “important” is written as **“IMPORTANT,”** both capitalized and bold, and is followed by text in italics. The italicized text is the message to be delivered.



Warnings

Warnings appear where overlooked details may cause damage to the equipment or result in personal injury. Warnings should be taken seriously. Warnings are easy to recognize. The word “warning” is written as **“WARNING,”** both capitalized and bold and is followed by text in italics. The italicized text is the warning message.



Cautions

Cautionary messages should also be heeded to help you reduce the chance of losing data or damaging the system. Cautions are easy to recognize. The word “caution” is written as **“CAUTION,”** both capitalized and bold and is followed by text in italics. The italicized text is the cautionary message.



Notes

These messages inform the reader of essential but non-critical information. These messages should be read carefully as any directions or instructions contained therein can help you avoid making mistakes. Notes are easy to recognize. The word “note” is written as **“NOTE,”** both capitalized and bold and is followed by text in italics. The italicized text is the cautionary message.

Galaxy® Aurora LS EOS / NumaRAID Updates

EOS/ NumaRAID is always looking to improve its abilities and therefore may upgrade the software.

Please contact your system vendor for the latest software updates.

NOTE that the version installed on your system should provide the complete functionality listed in the specification sheet/user's manual.

We provide special revisions for various application purposes. Therefore, DO NOT upgrade your software unless you fully understand what a revision will do. Problems that occur during the updating process may cause unrecoverable errors and system down time. Always consult technical personnel before proceeding with any firmware upgrade.

This page left blank intentionally

Section 1

Introduction and Overview

1.0 Introduction and Overview

1.1 Product Specifications

1.1.1 Overview

The Aurora LS model the newest member of the Galaxy family of RAID Storage System products. It is a (4U) rack mount solution that is designed for your ultra high speed data storage needs.



As with the earlier Galaxy® RAID products, the Galaxy® Aurora LS is characterized by many of the same outstanding features and attributes as those of other RAID family members. The most noticeable feature is that this RAID is blazingly fast while being surprisingly affordable. Other features include a preloaded Linux operating system and RAID Engine Software called EOS, or NumaRAID , which does all the work of a normal RAID controller without the cost and dependency of other ASIC based controllers. Of course speeds that exceed 1000Mbytes/ second would be no good without the host connectivity which is built into the unit. Aurora LS is capable of supporting 1/2/or 4 port 8Gb Fibre Channel HBA [host bus adapter] with connections to 1,2, 4 hosts directly or with SAN connectivity connect to many more. Optical cable connectivity is available in various lengths to make direct or SAN switch connections easy. Other features include, easy to use GUI storage management tools, integrated software functions that help ease configuration and use, ease of deployment in the network, as well as built-in tools to facilitate remote management and systems management. Our ultra quiet “dark space” design gives a very low noise signature. The innovative approach to fast RAID rebuild translates into a 100% of the rebuild occurring in 20% of the time a hardware based RAID would take. A major feature is only seen during the rebuild: no bandwidth performance loss.

1.1.2 Basic Features and Advantages

Galaxy® Aurora LS RAID products provide these important features and advantages:

Compact 4RU Steel and Aluminum Alloy enclosure with rack mount kit.

1000+ MB/s sustained bandwidth over 8Gb Fibre cables

Upgraded quad core XEON mother board

12 Drive SAS controller

64 bit Linux based OS

EOS embedded RAID Engine and GUI application

8Gb Fibre Channel SAN support

12 Removable Hot Swap Disk Drives

Multiple 2TB partition support for 32bit OS support

Web-based Graphical User Interface

Remote Maintenance with browser or command line

Remote Hardware Status monitoring

LUN Partitioning

Background Activities that include: RAID Rebuild; SMART condition polling; Enclosure and Media health monitoring and repair

- Failed drive , fan, and power supply indicators

Auto-rebuild while maintaining peak data bandwidth performance

Secured Administration Access

Multiple Management Network Interface Card (NIC) Support

Up to 12TB logical volume support

UPS Support and Network UPS Support

Console Tool as well as Remote Console

Email Alert Notification and Log Function

Supporting configurations that bridge to Fibre and Gbit Networks

1.2 Model Variations

1.2.1 Galaxy® Aurora LS Model Descriptions

The Aurora LS has 3 primary models with many storage variations:

GAURLS12-1FC8-18TA GALAXY AURORA LS 12BAY 4U TOWER / RACKMOUNT STORAGE APPLIANCE - 1X2.66GHZ CORE I7 CPU, 6GB (3X2GB) RAM, 12X1.5TB 7200RPM SATA DRIVES, LINUX OS & EOS APPLICATION ON DOM, SINGLE-PORT 8GBIT FC HBA, RAID 6, 1ST YEAR G1 WARRANTY

GAURLS12-2FC8-18TA GALAXY AURORA LS 12BAY 4U TOWER / RACKMOUNT STORAGE APPLIANCE - 1X2.66GHZ CORE I7 CPU, 6GB (3X2GB) RAM, 12X1.5TB 7200RPM SATA DRIVES, LINUX OS & EOS APPLICATION ON DOM, DUAL-PORT 8GBIT FC HBA, RAID 6, 1ST YEAR G1 WARRANTY

GAURLS12-4FC8-18TA GALAXY AURORA LS 12BAY 4U TOWER / RACKMOUNT STORAGE APPLIANCE - 1X2.66GHZ CORE I7 CPU, 6GB (3X2GB) RAM, 12X1.5TB 7200RPM SATA DRIVES, LINUX OS & EOS APPLICATION ON DOM, QUAD-PORT 8GBIT FC HBA, RAID 6, 1ST YEAR G1 WARRANTY

The Aurora LS models share the same basic setup, configuration, and administration so the main portion of the manual will discuss these functions.


For ease of purpose, the main portion of the manual will

be based on the GAURLS12-4FC8-18TA version of the Aurora LS.

1.3 Product Description

1.3.1 Description of Physical Components

See the figure below for a diagram of the front of the Galaxy Aurora LS.




The figure below shows a detailed diagram of the front controls area:



The figure on the following page shows a diagram of the rear of the Galaxy Aurora LS . Note that this configuration may be slightly different than your actual Aurora LS .

Figure 1.3.1c



- | | |
|--|--------------------------------|
| A) Upper Power Supply Module | S) Network Port 1 Link LED |
| B) Power Supply Reset PB | T) Network Port 2 Activity LED |
| C) Power Supply Overtemp LED | U) Network Port 2 |
| D) Fibre Channel HBA | V) Network Port 2 Link LED |
| E) Fibre Channel Port 1 | |
| F) 8Gb Fibre LED [both needed for 8Gb] | |
| G) SAS RAID Controller | |
| H) SAS Activity | |
| I) SAS Heartbeat | |
| K) PS/2 Mouse Connector | |
| L) PS/2 Keyboard Connector | |
| M) USB Ports | |
| N) Serial Port (Not used) | |
| O) Exhaust Fan Area | |
| P) VGA Connector | |
| Q) Network Port 1 Activity LED | |
| R) Network Port 1 | |

Facing the rear, the power supply module is located on the left. To the right of each power connector, is an over temp LED which is on if the power supply module is overheated because a fan is not operating and receiving power. If this LED goes on, it could mean that the power cable isn't operating properly, or there is a problem with the power supply, the module itself, or

the AC outlet. To remove a power supply module, you will have to gain access to the mounting screws inside the chassis. Contact Tech support for help with the removal procedure.

The two round connectors [K] [L] are for a PS/2 keyboard or a mouse. The green connector is for a mouse, and the purple connector is for a keyboard. To the right of these two connectors are USB connectors. These can be used for USB drive(s), memory key(s), hub(s), and/or a USB keyboard or mouse. To the right of the USB connectors is a green serial connector. It is not used. To the right of the serial connector is an analog VGA connector. You may attach a console monitor here. To the right of the VGA connector are two gigabit Ethernet ports. The left port (if you are facing the rear) is port 1, the right port is port 2.

The vertical slits on the right (called slots) hold the host adapters which are inside the system. Going from left to right, we see an empty slot, then the Fibre Channel host bus adapters [used to connect to your host system], with two LEDs associated with each connector; one LED indicates 4Gb host connectivity, two LEDs indicate 8Gb connectivity. Each Fibre connector can be connected to a unique host system. Next to the Fibre HBA are empty slots, followed by a SAS RAID Controller card that the 12 Disk drives are connected to. Be aware that a green LED blinks continuously indicating the processor on the adapter is functioning. The other blinks during activity.

1.3.2 Component specifications

The Aurora LS is a 4U 12-bay rack mountable storage enclosure that supports up to twelve hot-swappable hard disk drives. The latest Nehalem processor and Motherboard is the main hardware engine of the Aurora LS. This board supports:

- 2.66Ghz Core i7 CPU
- EOS RAID application and RAID GUI
- On board externally connected video, mouse, and keyboard
- On board dual 1Gb Ethernet ports
- Ships with 6GB DDR RAM
- Up to 3 PCI-X , 3 PCI-E slots
- Supports up to 12 x 3.5", 1.0" 3Gb SAS or SATA half-height hard disk drives [storage size and speeds vary depending on model]
- Twelve hot-swappable hard disk drive bays
- Integrated backplane design that supports 3Gb SATA / SAS Disk Interface
- Built-in environment controller
- Enclosure management controller
- Ultra quiet 'dark space' fan design
- Advanced thermal design with hot-swappable fans
- Front panel LED Alarm and Function indicators
- Shock and vibration proof design for high reliability
- Dimensions: 13.1x 44.65x 56.1 cm (7.0 x17.2 x 26.1in)
- Weight: Gross weight (including carton): 27.5kg (58.7 lbs) with drives, 33.1 kg (72.0 lbs) with 12 drives
- Power Supply: 865 Watt, 100-240 Vac auto-ranging, 50-60 Hz
- Operational Wattage / Amps: 275 watts / 2.3 amps

- Ventilation: Quiet 'dark space' 3 fan design
- Environment Controller Internal Temperature - visible and audio alarm

1.3.3 RAID storage specifications

The Aurora LS has a sophisticated built in RAID software and drives that are pre-configured and prepared for you so it would be plug and play for most users. By default, the Aurora LS RAID has been configured into one RAID 6 logical volume. For 32 bit Windows XP configurations, multiple 2TB volumes would have been created for you.

RAID 6 with its dual parity drive protection has been found to be the most protective and least costly way of guarding against not only initially failed SATA disk drives but primarily against the total loss of the RAID data because a second SATA drive detects an error during the RAID rebuild process. A RAID 5 configuration in that scenario would cause the RAID not to rebuild properly.

1.3.4 Embedded OS features



ⓘ Important: The Aurora LS EULA restricts you, the user, from loading any other software, such as application software, onto the Aurora LS. Tampering with, loading or using any other software voids the license agreement.

Each Aurora LS is preloaded at the factory with its base operating system, RAID application, installation, administration and optional SAN software. The code is loaded onto the system's boot drives.

In addition to the operating system and basic EOS embedded application software, each unit contains a web based browser interface which simplifies remote configuration and administration tasks.

Specifically, the units come preconfigured with the following functions:

EOS: Linux based RAID application and User configuration / troubleshooting interface

Remote system administration:

Administrative tasks can be performed in the Web-based GUI

Alternate administrative task performed using Windows Terminal Service

Advanced management functions available via Windows Terminal Service

Optional SAN Management Software

1.4 Mounting / Securing Aurora LS

1.4.1 Rack Mounting the Aurora LS

The Aurora LS is a rack mounted chassis. Mounting holes on the front panel are set to RETMA spacing and will fit into any standard 19" equipment rack.

Rack Equipment Precautions

These precautions and directions should be used only as an information source for planning your Aurora LS deployment. Avoid personal injury and equipment damage by following accepted safety practices.

Floor Loading

 **CAUTION:** Ensure proper floor support and ensure that the floor loading specifications are adhered to. Failure to do so may result in physical injury or damage to the equipment and the facility.

Deployment of rack servers, related equipment, and cables exceeds 1800 pounds for a single 42U rack.

External cable weight contributes to overall weight of the rack installation. Carefully consider cable weight in all designs

Installation Requirement

 **CAUTION:** Be aware of the center of gravity and tipping hazards. Installation should be such that a hazardous stability condition is avoided due to uneven loading. We recommend that the rack footings extend 10 inches from the front and back of any rack equipments 22U or higher.

Adequate stabilization measures are required. Ensure that the entire rack assembly is properly secured and that all personnel are trained in proper maintenance and operation procedures. Tipping hazards include personal injury and death.

Power Input and Grounding

 **CAUTION:** Ensure your installation has adequate power supply and branch circuit protection.

Check nameplate ratings to assure there is no overloading of supply circuits that could have an effect on over current protection and supply wiring. Reliable grounding of this equipment must be maintained. Particular attention should be given to supply connections when connecting to power strips, rather than direct connections to the branch circuit.

Thermal Dissipation Requirement

 **CAUTION:** Thermal dissipation requirements of this equipment deployment mandate minimum unrestricted airspace of three inches in both the front and the rear. The ambient within the rack may be greater than room ambient. Installation should be such that the amount of air flow required for safe operation is not compromised. The maximum temperature for the equipment in this environment is 122°F (50°C). Consideration should be given to the maximum rated ambient.

1.4.2 Installation Sequence

 **CAUTION** It is strongly recommended to securely fasten the mounting rack to the floor or wall to eliminate any possibility of tipping of the rack. This is especially important if you decide to install several Aurora LS 'chassis' in the top of the rack.

A brief overview of Aurora LS installation follows:

1. Select an appropriate site for the rack.
2. Unpack the Aurora LS and rack mounting hardware.
3. Attach the rack mounting hardware to the rack and to the Aurora LS.

4. Mount the Aurora LS into the rack.

5. Connect the cables.

Decide on an appropriate location for the Galaxy Aurora LS . It is best if the unit is kept away from heat or from where high electromagnetic fields that may exist. If you are installing the unit into a rack, make sure the rack is in the proper location prior to installation. Moving the Galaxy Aurora LS while it is installed into the rack is not recommended.

The Galaxy Aurora LS , requires 4 rack units of vertical clearance (7 inches), and a depth of 28 inches. It is recommended that you mount it in a rack which is at least 30 inches deep.

Airflow for the unit comes in through right side and the front. Heat exhaust is from the rear of the unit. It is important that airflow at the front or the rear not be blocked.

The rack slides permit the unit to slide out of the front of the rack. There are latches on the sides of the slides, and if you are planning on removing the unit from the rack to service or transport it, sufficient clearance should be available to allow you to activate the latches and unlatch the slides.

If the rack is on wheels, be sure to use the wheel locks when installing or removing the Galaxy Aurora LS from the rack. If the rack does not have wheel locks, place something against the wheels to prevent movement, or if your rack is equipped with leveling jacks, extend the jacks to make sure the rack stays level during installation. Always make sure the rack is completely immobile before installing or removing any components. Never extend more than one component from the rack at the same time.

There is a set of slides included with the Galaxy Aurora LS . The slides are required for rack-mounting the unit, and the slides must be mounted with the rear extensions installed into the rack. The weight of the unit is sufficient that if this were not performed, damage would result to the unit, the slides, or the rack if installed.


When installing the slides, loosely attach the rear end of the slide to the front end, then screw the front and rear rack portions of the slides into the rack. Finally, tighten the screws between the two ends. Repeat this process for the other side. Once the slides are installed in the rack, slide the unit into the slides. Slide extensions are included in case the rack is deeper.

1.4.2.1 Ball Bearing Slide Rail Rack Installation

Unpack the package box and locate the materials and documentation necessary for rack mounting. All the equipment needed to install the server into the rack cabinet is included.

Follow the instructions for each of these illustrations


Kit Contents: the rack mounting kit include:




- ❶ Rack mounting brackets (4)
- ❷ Slide rail assembly (2)
- ❸ Bags of screws and nuts (8 x hex nut #8-32, 8 x screw #8-32 1/4 inch, 16x screw #10-32-1/2 inch)

Installing the Rack Rails


Determine where you want to place the Aurora LS in the rack. Position the fixed rack rail/sliding rail guide assemblies at the desired location in the rack, keeping the sliding rail guide facing the inside of the rack. Screw the assembly securely to the rack using the brackets provided. Attach the other assembly to the other side of the rack, making sure both are at the exact same height and with the rail guides facing inward.



Remove the Server Rail




Fasten Bracket and Outer Piece Rail Together



- 1 Align the mounting bracket with the outer rail
- 2 Fasten but not tighten the Bracket with the screws provided

Mount Bracket and Outer Rail to the Rack Cabinet



Installing the System into the Rack

You should now have rails attached to both the chassis and the rack unit. The next step is to install the system into the rack. You should have two brackets in the rack mount kit. Install these first keeping in mind that they are left/right specific (marked with "L" and "R"). Then, line up the rear of the chassis rails with the front of the rack rails. Slide the chassis rails into the rack rails, keeping the pressure even on both sides (you may have to depress the locking tabs when inserting).

When the system has been pushed completely into the rack, you should hear the locking tabs "click". Finish by inserting and tightening the thumbscrews that hold the front of the chassis to the rack. This completes the installation and rack mounting process.

 **CAUTION** Due to the weight of the chassis with the peripherals installed, lifting the chassis and attaching it to the cabinet may need additional manpower. If needed, use an appropriate lifting device.

Section 2

Basic Setup

2.0 Basic Setup

2.1 Drive integration and Cable Connections

2.1.1 Indicators and switch descriptions Figure 2.1

The Aurora LS front panel has indicators for good and fault conditions and activity. Green LEDs indicate good condition, red LEDs indicate a problem that will also log an error and send an email. The alarm reset needs to be depressed to silence the alarm. The Reset PB is used to restart the Aurora LS. The Power PB is used to power up the Aurora LS.

Figure 2.1



The power switch is used to turn the unit on. However, do not use it to turn the unit off, unless there is no other way. To turn on the unit, press the power switch momentarily. To turn it off, press and hold it for 8 seconds. The reset switch also should not be used unless there is no

alternative. To the right of the two switches is the Power LED.  This illuminates when

power is on. To the right of the power LED is the Aurora LS OS DOM activity LED . This LED will light intermittently during normal operation. Next to the power LED are two network

LEDs.   These LEDs will light when there is activity from the management network ports they correspond to on the rear. Next to these is a temperature / fan fail warning

LED  . If the temperature inside the system becomes too high, this LED will illuminate.

Next to the temperature warning LED is a power fail LED  . If there is something wrong with the power supply fan , this LED will illuminate. The USB2 ports are active and should only be used with a USB based keyboard and mouse.

2.1.2 Installing drives into the Aurora LS Figure 2.2


The Galaxy Aurora LS features 12 removable drives. They have been shipped separately to insure the Aurora LS would not incur shipping damages from a possible shipping related shock to the drives or backplane.

 **CAUTION:** Be aware that the Aurora LS's file system does not support drive roaming. Drives must be installed and must be placed into their prepared slots for the RAID set to operate properly.

The drives will be tagged with numbers 1-12. Place them in their assigned numbered slot in the Aurora LS chassis as shown below.

NOTE THAT DRIVE SLOT 13 HAS AN EMPTY DRIVE CANISTER!

Figure 2.2



The drives are simple to install. Simply unwrap and push each drive into each empty drive opening as far as it will go, then push the handle in until the red button clicks into place. Each of the drive modules in the Galaxy Aurora LS has two LEDs the upper LED flashes for disk activity, while the lower LED is used for errors and flash ID use. The RAID's EOS software will automatically find all drives. To remove a drive module, push the red button until the black handle pops out. Then pull the handle until it is sticking straight forward, and carefully pull the drive out by the handle. To reinstall a drive, make sure the handle is sticking out of the module (if it's not, push the red button to release the handle),


The Aurora LS OS has been preloaded and RAID storage preconfigured to be ready for you to power up and start configuring it for use. Before powering up, make the cable connections to , ethernet, power, keyboard, and monitor [in certain cases these components and cables are provided].

2.1.3 Connecting Cables Figure 2.3

See the illustration for the cable locations and connectivity.

For safety reasons we recommend the cables be connected in the following order:

Connect one power cord to an active powered AC outlet, then connect the other end to the rear of the Galaxy Aurora LS. You will hear a fan get loud, then get quiet – this is normal and nothing to be alarmed about.



Then connect the Ethernet, Fibre Channel, monitor, keyboard and mouse.

The 8Gb Fibre Channel connection can either be connected point-to-point (i.e. directly to another computer with 4Gb or 8Gb Fibre Channel host adapter), or can be connected to a 4Gb or 8Gb Fibre Channel switch.

When all cables are installed, one or more of the Ethernet activity LEDs on the front of the unit may blink.

Power up the Galaxy Aurora LS by momentarily pushing the front Power switch. The Galaxy Aurora LS, will take several minutes to boot.


Once powered up and all indicators are green, the Aurora LS is online to your system and should be seen as a typical RAID storage system. For further use of the GUI for sense, setting, and configuration information, continue with the following information.

2.2 Configuration Setup


2.2.1 Setting up Ethernet Connectivity on a Windows Client

For you to administer Aurora LS, setup remote maintenance, or proceed with SAN usage you need to be able to see the Aurora LS with a standard internet browser over ethernet from your client. The process below will allow the client to talk to the Aurora LS over ethernet on a Windows Client. Contact your Network Administrator for support.

Proceed to the TCP/IP settings area of your particular client station, ie Windows control panel network settings and select properties. Select the TCP/IP listing and click properties:



Click the button to 'Use the following IP address':



Setup the

IP address to :192.168.1.2

Subnet mask to : 255.255.255.0

Default gateway to: blank

DNS server info should be blank.

Click OK and your client can now see the Aurora LS over Ethernet using as standard Internet Browser.



The Aurora LS has been setup with a fixed default IP address of :
192.168.1.129

2.2.2 Installing Fibre Channel HBA and drivers on Aurora LS Clients

Consult with your local Aurora LS reseller for Windows, Linux, and Apple client HBA information.

Go to the various Linux, Windows, or Apple File system preparation section of this manual to prepare the Aurora LS LUN for your clients.

2.2.3 Linux Client RAID Connections and LUN Preparation

After Fibre channel HBA drivers are installed and loaded (which is not covered in this manual), you should already have the block device representing the LUN mounted. If you type the following command you should get a list of mounted storage LUNs:

```
lsscsi[enter]
```

the following response will be displayed:

```
[0:0:0:0] disk ATA HDS722516VLAT80 V34O /dev/sda
[0:0:1:0] cd/dvd PIONEER DVD-RW DVR-109 1.40 /dev/sr0
[2:0:0:0] disk GalaxyIB MyLUN 2091 /dev/sdb
```

In the example above, the last line shows the Aurora LS LUN [GalaxyIB My LUN]. The Aurora LS device manufacturer is shown as GalaxyIB, with the My LUN name as the model name. The version number, 2091, is the version of the Aurora LS driver. Finally, you are most interested in the device name on the right [/dev/sdb]. The next step for preparing to use this LUN is to label the device, and create a partition on it. This is done with the Linux 'parted' command, by typing the following:



CAUTION: This procedure erases all data on the LUN.



Important : Be very careful typing these keyed entries in **bold type**.

Go to a new prompt and enter:

```
parted /dev/sdb[enter] the responding command line interface is displayed as:
```

```
GNU Parted 1.8.7
```

```
Using /dev/sdb
```

```
Welcome to GNU Parted! Type 'help' to view a list of commands.
```

```
(parted) mklabel[enter]
```

```
Warning: The existing disk label on /dev/sdb will be destroyed and all data on
this disk will be lost. Do you want to continue?
```

```
Yes/No? Yes[enter]
```

```
New disk label type? [gpt]? gpt[enter]
```

```
(parted) mkpart[enter]
```

```
Partition name? []? mypart[enter]
```

```
File system type? [ext2]? ext3[enter]
```

```
Start? 0[enter]
```

```
End? -1[enter]
```

```
(parted) quit[enter]
```

In the example above, the /dev/sdb typed after the parted command specifies the device to partition as seen from the lsscsi command. When entering the make a label command

[mklable], it gives a warning about an existing label – you may or may not get this warning – this is not an error. A label is basically a data element which is written to the device on its outer-most sector, which describes very generally how it is going to be used. The main options are mbr and gpt. mbr is for devices which are 2TB in capacity or less. gpt is for any size device – it can also be used for devices which are 2TB in capacity or less. When creating the partition, the name “mypart” was given. The partition name really isn’t used outside of parted itself, so it doesn’t really matter what you name it, but it does have to have a name, preferably unique. Also the file system chosen for this example was ext3. Other file systems may be used on your client – some offer features that others do not have and vice-versa. Because this is showing up as a block device on the client, the array itself doesn’t have to support the file system being used. The ‘Start?’ entry of ‘0’ indicates the starting sector number is 0. The ‘End?’ entry of “-1” indicates that the end of the partition is on the last sector. It’s possible to have multiple partitions, but for this example, the entire LUN is used. Consult with tech support for partition size options. In this case you have created partition 1 but still need to create a file system on it.

The file system has to be created on that partition. The device in the example is /dev/sdb, however the partition is specified by typing the partition number after the device – in this case /dev/sdb1. In the example, the ext3 file system was specified. The command to create the file system has to match the file system selected during ‘parted’. To create the ext3 file system now on partition /dev/sdb1, ‘make file system’ [mkfs] command is used. type the following:

```
mkfs.ext3 /dev/sdb1[enter]
mke2fs 1.40.2 (12-Jul-2007)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
131072000 inodes, 262143991 blocks
13107199 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=4294967296
8000 block groups
32768 blocks per group, 32768 fragments per group,16384 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624, 11239424, 20480000, 23887872, 71663616, 78675968,
    102400000, 214990848
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
This filesystem will be automatically checked every 27 mounts or
180 days, whichever comes first. Use tune2fs -c or -i to override.
```

The amount of time it takes to create the file system will vary, depending on the file system chosen, the LUN capacity, the drive speeds, connection type, etc. Some file systems create in just seconds, while others can take minutes or hours. In the example above the ext3 file system creation took approximately 2 minutes.

The partition is prepared but must be mounted to use the LUN by the Linux clients. Here's the command:

```
mount /dev/sdb1 /mnt[enter]
```

In this example, you are mounting the ext3 partition /dev/sdb1, to a pre-existing folder, /mnt. You can create your own mount points, by using the following commands:

```
mkdir {/folderpath}[enter]
```

```
chmod 777 {/folderpath}[enter]
```

For example, to mount the array to /root/bob, you would type the following:

```
mkdir /root/bob[enter]
```

```
chmod 777 /root/bob[enter]
```


```
mount /dev/sdb1 /root/bob[enter]
```

Once the mount point is created, it doesn't have to be recreated each time –just use the mount command.


Your Aurora LS LUN is now available for use by Linux clients.

2.2.4 Windows Client RAID Connections and LUN Preparation


After the Fibre channel HBA drivers are installed and loaded (which is not covered in this manual), system rebooted and cabled to the array, begin by left-clicking on the [Windows logo](#) (Or Start Menu) in the lower left corner of the screen: Note that the instructions here are for Vista Ultimate/64 but other versions are similar.




This will cause the start menu to pop up. Your screen will look different – not every computer has the same programs in the list. Along the right side of the menu is a grey area (in the image above), move the mouse pointer to [Computer](#) and right-click on it:



This will launch an additional menu. Left-click on [Manage](#) on this new menu:



The Computer Management window will open - On the left side of the screen, left-click on Disk Management under Storage. If it is not visible, either turn down the arrow to the left of Storage or scroll down to it:



The right side of the screen will change. If this is the first time that this LUN has been formatted for Windows, an Initialize Disk popup will appear on top of the disk management window. This warning will usually also only appear on 64-bit OSes. If you are running a 32-bit OS, and your LUN is greater than 2TB, it won't show up at all in disk management, because Windows 32-bit OSes have a 2TB physical device size limit.

Important :The Aurora LS does have the ability to create larger than 2TB LUNs for 32-bit Windows but the GUI LUN creation method needs to be used in Section 3.




CAUTION: At this point, the LUN will be relabeled from the client – it may erase any data that was on the LUN.


Left-click on the bubble next to GPT. Then left-click on the OK button:



The Disk Management window will open. In the example below, a 1TB LUN was used – it is appearing as Disk 1. To the right of Disk 1, a large rectangle with a black bar running across the top. Right-click in the white rectangular area just below the black bar:




The following pop-up menu will appear , left-click on New Simple Volume...




This will open the New Simple Volume Wizard. Left-click on the Next > button to continue:



The 'Specify volume size window will open. Use the default values. Left-click on the Next > button to continue:




The 'assign drive letter' window opens. Use the default and note the letter. Click on the Next > button to continue:



The 'format partition' window opens. Leave all values at default except Volume label. Left click the volume label and enter a preferred name for the partition:




On the same window left-click on the “Perform a quick format,” checkbox (So it is checked) then left-click on the Next > button:



The ‘completing the simple volume..’ window opens. This is the final window of the wizard, which shows all of the settings that were selected and provides the last chance to go back and make any changes before the LUN is formatted and volume created on it. If everything looks OK, click on the Finish button to continue:




When the partitioning is finished, the New Simple Volume Wizard will close, and you will be returned to the Disk Management screen. After a few moments (less than a minute), the Disk Management screen will update the information about the new volume as follows, and your volume is ready to use:



2.2.5 Apple OSX Client RAID Connections and LUN Preparation

Refer to the Fibre channel HBA installation instructions to install your HBA and drivers into your Apple OSX clients. This document also uses OS/X 10.5 as an example – all versions of OS/X supported by the Fibre Channel host adapter should work and have almost identical setup procedures (From 10.2 to 10.6, except where noted). Once you have installed your host adapter, connected the fibre cable, and rebooted, you may see the following popup window. If you get this warning, it will save all of the steps necessary in setting up the Aurora LS with Apple Disk Utility. So if the Disk Insertion warning does appear, click on the Initialize... button:




Initializing the Aurora LS is the purpose of this procedure so if this popup did not come up, or if you closed it by accident, or if it closed by itself, or if you want to know how to get into the Apple Disk Utility and setup the initialize manually, follow these steps:




On your desktop, you will see an icon **Macintosh HD** (usually in the upper-right corner of the screen), which represents your boot drive. Double-click this icon to open your boot drive.


The Finder will open. If you have not seen or used the finder before, contact Tech Support for assistance. Click on Applications, which is near the top of the list:




The next column to the right will populate, showing the contents of the Applications folder. On most systems, this new column will be too large to fit on the screen, so you will need to scroll all the way to the bottom. Click on the slider, and drag it down and navigate and click on the Utilities folder:




The next column to the right will populate, showing the contents of the Utilities folder, double-click on Disk Utility:




Apple Disk Utility will open - You will see the LUN listed on the left – in the example above, it is a 1TB LUN, showing GalaxyIB testlun1 Media. Click on the LUN to select it:




On the upper right is a series of tabs. Click on the Partition tab to select it if it is not already selected:




In the middle of the screen, click on Current: in the “Volume Scheme” pulldown to expose a partition list:




Drag down to set the number of partitions to 1 Partition, then release the mouse button:



Click in the white text area next to Name:, and type a name for the volume:



At the bottom right, click on the Apply button:



A popup warning will appear.




CAUTION: Proceeding beyond this point will erase the LUN.

Click on the Partition button:



The partition and volume creation process will begin – this will only take a few seconds. When it is done, if you have OS/X 10.5 or above, another popup window will appear about Time Machine: Click on the Cancel button:



Apple Disk Utility as follows – the process is complete and the volume appears on the desktop:

2.3 Remote Administration

2.3.1 Using a Browser and Logging into the Aurora LS

The Galaxy Aurora LS is managed by a browser or command line interface. For ease of use the user should use a browser remotely to verify the basic operations and functionality.

This is accessed by opening a browser, and typing the following URL:


<http://192.168.1.129:10000>

You will see a login window.

The login is: admin

It is case-sensitive, and the password is: password

When you log in via the GUI, you will see the Galaxy Aurora LS Home Screen, which looks like this:



System hostname	GalaxyIB
Operating system	Redhat Enterprise Linux 5
Webmin version	1.441
Time on system	Thu Sep 24 10:34:47 2009
Kernel and CPU	Linux 2.6.18-92.el5 on x86_64
CPU load averages	0.16 (1 min) 0.25 (5 mins) 0.17 (15 mins)
Real memory	7.67 GB total, 3.06 GB used
Virtual memory	767.99 MB total, 0 bytes used
Local disk space	6.51 GB total, 2.51 GB used

Discussion of Managing the Aurora LS follows


Section 3

Aurora LS Management

3.0 Aurora LS GUI Detailed Operations

The GUI Menu provides you with simple and basic functions that can give you the overall status of the Aurora LS. Once logged in through a browser [<http://192.168.1.129:10000>] the following functions and features are available to the client.


3.1.0 GUI Menu Details and Functions



3.1.1 Main GUI screen page details and Quick Start functions

On your local computer you enter the GUI through the Mozilla Firefox web browser. Once inside the browser, enter the following URL: <http://192.168.1.129:10000>. This will give you a login prompt. The user name is **admin**, with the password being **password**.


The initial web admin page opens. In the Webmin menu on the left, expand the selection called “Hardware.” Below this click on NumaRAID GUI – this will launch the Main GUI Screen as follows:



System information displayed on the Webmin page:

System hostname	NumaRAID
Operating system	Redhat Enterprise Linux 5
Webmin version	1.441
Time on system	Thu Sep 24 10:34:47 2009
Kernel and CPU	Linux 2.6.18-92.el5 on x86_64
CPU load averages	0.16 (1 min) 0.25 (5 mins) 0.17 (15 mins)
Real memory	7.67 GB total, 3.06 GB used
Virtual memory	767.99 MB total, 0 bytes used
Local disk space	6.51 GB total, 2.51 GB used

The group will expand and will show an item below it called NumaRAID GUI. click on the NumaRAID GUI item under the hardware group to launch the main GUI page.



Hardware menu items:

- CD Burner
- GRUB Boot Loader
- Linux RAID
- Logical Volume Management
- NumaRAID GUI**
- Partitions on Local Disks
- Printer Administration
- SMART Drive Status
- System Time

Main GUI Screen:

Module Config **Left-Click here**

RAID Status							
RAID Name	Cache Size(M)	RAID Level	First Slot	Device Count	RAID Size(G)	Code Rev	Status
No RAID(s) defined							

Configuration Status	
	Type
Details	Configurations and Log Files
Details	Users
Details	Parameters
Details	Data Rate Statistics

RAID Enclosure Status		
	Type	Status
Details	Slots	OK
Details	Sensors	OK
Details	Adapters	
Power Control	Power Off ▾	

On the upper left is a link called Module Config - this is used to enable or disable the ability to change settings on the other screens which allow changes. The Aurora LS's main GUI "NumaRAID GUI Main Page," version number for the GUI (In this case 1.2.10). Below these are a series of three tables.

The first table shows the RAID Status. A RAID is a set of slots or disks, set to act in conjunction as one larger device. A RAID does not necessarily need to contain all of the disks in the array. Because of this, there are three possible things you could see in this table: If no RAID(s) are defined, it will say "No Raids defined", followed by the option to create a RAID. If RAID(s) are defined, but drives are still available, you will see a list of the RAID(s), with Details buttons next to each of them, followed by the option to create a RAID. If all of the drives are used in RAID(s) (as in the example above), you will see the list of RAID(s), but will not be given the option to create any new RAID(s)..

Click on Module Config – On the top of the Main GUI screen.

Configuration
For module NumaRAID GUI

Configurable options for NumaRAID GUI		
Allow Configuration Update	☐ No ☒ Yes	Left-Click Here then Here
Allow Parameter Update	☐ No ☒ Yes	
Save		
Return to index		

Click the yes buttons and click save. Return to the Main Screen which now displays the information about the RAID.

3.1.2 RAID Creation, Status, and other RAID configuration information

Although you have RAID created already you will need to know how to create a RAID (In this case, the example used is when no RAID exists):

RAID Status

RAID Name	Cache Size(M)	RAID Level	First Slot	Device Count	RAID Size(G)	Code Rev	Status
No RAID's defined							
Create		1GB	RAID6	0	16		

Do not click the Create button until everything else on the row is set correctly. To the right of the create button is where you give the RAID a unique name. The RAID requires a unique name, because it is referenced in a lot of places within Aurora LS, and would not be easy to identify if there was more than one RAID with the same name.

The next setting is the cache size. Cache is a designated part of the RAM in the array, used to hold data while waiting to go to the drives, or coming from the drives, waiting for the host. It is used to increase speed, because compared to the speed of the RAM, the speed of the drives are relatively slow, and the speed going to the host computer itself is unpredictable.

Important: The cache size selected is directly subtracted from the RAM in the array, so care must be taken so that not all of the RAM is not used up. For example, if you have 6GB of RAM, and already have an RAID defined which has a cache size of 4GB, then you don't have enough free RAM to create another RAID. Also, assume the operating system of the array itself takes about 2GB of RAM. In general, a larger cache yields greater performance. Once you know what cache size you would like to use, select it by left-clicking on the down arrow under Cache Size, then scroll down to the size that you would like, and left-click on it.

The next setting is the RAID level – it can be RAID 0 or RAID 6. With RAID 0, you get the capacity and potential speed of all of the disks, however if a single drive fails, you will lose access to all of your data. With RAID 6, you lose capacity equivalent to two of the drives, and get nearly the same speed, however up to two drives can fail and your data will still be accessible and at full speed.

The next setting is the number of the first slot/device to use for the RAID. Use the device count, where you would select the number of slots/drives to use in the RAID. The numbers used for the starting slot and device count must be contiguous - for example, if you specified that the starting slot was 5, and a device count of 4, slots 5, 6, 7, and 8 must be available.

Once you have made these selections, then you would left-click on the Create button.

For example, consider these settings:

RAID Status

RAID Name	Cache Size(M)	RAID Level	First Slot	Device Count	RAID Size(G)	Code Rev	Status
No RAID's defined							
Create	SubZero	4GB	RAID6	0	16		

In this example, "SubZero" was chosen for the name of the RAID, and a cache size of 4GB was selected. It is set to be RAID6 (Was already selected by default), and the RAID is set to use 16 devices starting with drive/slot 0. When the Create button was clicked, it indicated the command completed successfully. The process returned to the Main GUI Screen, here's how the RAID Status table looked:

RAID Status

RAID Name	Cache Size(M)	RAID Level	First Slot	Device Count	RAID Size(G)	Code Rev	Status
Details	SubZero	4000	6	0	16	1914	2089
							Online

You can see the RAID Name was set to SubZero, the Cache size was set to 4000 Megabytes (4GB), the RAID Level was set to RAID 6, the First Slot (Starting drive number) was set to 0, the number of devices was set to 16, and the RAID Size (The total usable capacity of the RAID in Gigabytes), in this case 1914GB or 1.9TB. The Code Rev is the version of the driver that is currently on the array – in this example, 2089. The Status shows whether the RAID is currently online or offline (in this case, online). Also notice that the Create option is no longer available, because all of the slots were used to create the RAID.

3.1.3 RAID Details

RAID Details

	RAID Name	Cache Size(M)	Cache Stripes	RAID Size(G)	RAID Level	Device Count	Status
Delete	BigFoot	1000	618	1093	0	8	Online
Scan/See Performance Stats							

No LUNs defined

LUN Creation

	LUN Name	Size(G)(Optional)	Offset(G)(Optional)
Create			

RAID Drive Details by Slot

Slot	Vendor	Model	Rev	Capacity(G)	Device	Linux Device	Status
0	SEAGATE	ST3146855SS	0002	146	scsi-35000c5000696adf3	sdb	Healthy
1	SEAGATE	ST3146855SS	0002	146	scsi-35000c50007bd20ef	sde	Healthy
2	SEAGATE	ST3146855SS	0002	146	scsi-35000c5000696997f	sdh	Healthy
3	SEAGATE	ST3146855SS	0002	146	scsi-35000c5000696999f	sdh	Healthy
4	SEAGATE	ST3146855SS	0002	146	scsi-35000c50002f2e48f	sdh	Healthy
5	SEAGATE	ST3146855SS	0002	146	scsi-35000c50006965a83	sdf	Healthy
6	SEAGATE	ST3146855SS	0002	146	scsi-35000c5000696ae97	sdi	Healthy
7	SEAGATE	ST3146855SS	0002	146	scsi-35000c500069698fb	sdi	Healthy

[Return to NumaRAID GUI Main Page](#)

The RAID details screen is used to view information about the devices which make up a RAID, as well as view and create LUN(s) on the RAID, and test the RAID/LUN/Drives.

At the top, we see the status of the RAID, very similar to the main screen - it shows the name of the RAID (in this example, Bigfoot), the cache size in Megabytes (In this example, 1000 Megabytes or 1 Gigabyte), the number of cache stripes (in this example 618). The way cache stripes are used is the stripe size (the default is 128KB) x the number of drives x the cache stripes is the amount of RAM of the cache that is used only for data caching. The columns to the right of cache stripes show the total capacity of the RAID (in Gigabytes – in this example, 1093 Gigabytes), the RAID Level (0 or 6), the number of devices which make up the RAID, and the overall status of the RAID. On the left is a [Delete](#) button – this is used to delete a RAID, however a RAID can not be deleted unless no LUN(s) exist on that RAID.

At the bottom of the RAID status table is a [Scan/See Performance Stats](#) button which takes you to a screen where you can scan and see performance statistics for the RAID. This will be covered later.

Below the RAID status, is a table of LUN(s), if any. A LUN is a logical portion of a RAID, which is presented to a client system as a block device. It is logical, because a LUN only exists in the configuration - Nothing is written to the data area of a RAID to define a LUN. In my example above, there are no LUN(s) defined yet, so the table just says "No Luns Defined." At least one LUN must exist in order for the array to be seen by a client.

Below the LUN table is an area where you can create a LUN. All entries must be made before left-clicking on the [Create](#) button. To the right of the create button is an area where you can enter the LUN name - all LUN(s) should have unique names. By default, with no size or offset, if the LUN is created, it will be the full size of the RAID that you are creating it on, otherwise the size entered is the size of the LUN (in Gigabytes), and the offset is where to start the LUN (in Gigabytes). The area encompassed by a LUN can not be used by another LUN, and it must be contiguous. For example, if you had a RAID that was 8TB, with (4) 2TB LUNs on it, then if you deleted LUNs 1 and 3, you could not create a 4TB LUN in that space, because LUN 2 would be in the way. Here is an example, where a single LUN was created, called MyLun. All that was done to create this lun was "MyLun" was typed for the name, then the

Create button was clicked. Back on the Raid Details screen, here is how the LUN status now appears:

Lun Status

	Lun Name	RAID Name	Size(G)	Offset(G)
Details	MyLun	BigFoot	1093	0

It shows the name of the LUN as MyLun, then it shows the name of the RAID that it belongs to (BigFoot), the size/capacity of the LUN (in Gigabytes – in this example, 1093 Gigabytes), and the offset (Starting point – also in Gigabytes – in this case, 0). The Details button launches Lun Details, where Initiator and Target assignments are performed. This is covered in more detail later.

Below the LUN creation area of the RAID Details screen is the RAID Drive Details by Slot table:

RAID Drive Details by Slot

Slot	Vendor	Model	Rev	Capacity(G)	Device	Linux Device	Status
0	SEAGATE	ST3146855SS	0002	146	scsi-35000c5000696adf3	sdb	Healthy
1	SEAGATE	ST3146855SS	0002	146	scsi-35000c50007bd20ef	sde	Healthy
2	SEAGATE	ST3146855SS	0002	146	scsi-35000c5000696997f	sdh	Healthy
3	SEAGATE	ST3146855SS	0002	146	scsi-35000c5000696991f	sdh	Healthy
4	SEAGATE	ST3146855SS	0002	146	scsi-35000c50002f2e48f	sdh	Healthy
5	SEAGATE	ST3146855SS	0002	146	scsi-35000c50006965a83	sdf	Healthy
6	SEAGATE	ST3146855SS	0002	146	scsi-35000c5000696ae97	sdi	Healthy
7	SEAGATE	ST3146855SS	0002	146	scsi-35000c500069698fb	sdi	Healthy

This table shows (in slot order), the slot number, each drive with the manufacturer, the model, the firmware version, the capacity (in GB), the Linux by-id device name, the Linux short device name, and the status. In the device column, there is an important distinction, depending on whether SAS or SATA drives are used. With SAS drives, the hexadecimal number after “scsi-” is the SAS address of the drive (The SAS addresses are printed on the drives). On SATA drives, the last 8 characters of this device name will be the serial number of the drive.

At the bottom of the RAID Details screen, you may left-click on the [Return to NumaRAID GUI Main Page](#) link if you wish to return to the Main GUI Screen.

3.1.4 Scan / Performance Results

RAID Details

RAID Name	Cache Size(M)	Cache Stripes	RAID Size(G)	RAID Level	Device Count	Status
LittleFootOne	1000	1428	546	0	4	Online

RAID Surface Scan

RAID Name	Type	Size(G)	Offset
LittleFootOne	Read Scan	1G	0%

Real Time Response Characteristics - In Milliseconds - Sequential Read

Drive	0-15	16-31	32-47	48-63	64-79	80-95	96-111	112-127
scsi-35000c50006969a43	1	0	0	0	0	0	0	0
scsi-35000c50007bda8f3	1	0	0	0	0	0	0	0
scsi-35000c50006965433	1	0	0	0	0	0	0	0
scsi-35000c50007bd5297	1	0	0	0	0	0	0	0

Real Time Response Characteristics - In Milliseconds - Sequential Write

Drive	0-15	16-31	32-47	48-63	64-79	80-95	96-111	112-127
scsi-35000c50006969a43	1	0	0	0	0	0	0	0
scsi-35000c50007bda8f3	1	0	0	0	0	0	0	0
scsi-35000c50006965433	1	0	0	0	0	0	0	0
scsi-35000c50007bd5297	1	0	0	0	0	0	0	0

[Reset Performance Response Counters](#)

[Return to NumaRAID GUI Main Page](#)

When you click the 'Scan / See Performance Stats' button on the RAID Details page, the performance page opens as the example above shows.

This is a very important screen which can help troubleshoot problematic hard drives: At the top, the RAID Details table shows the name of the selected RAID, the cache size (in Megabytes), the number of cache stripes, the RAID size (in Gigabytes), the RAID Level (0 or 6), the device count (the number of devices which make up the RAID), and the overall RAID status. RAID Surface scan will be discussed later.

Real Time Response times are displayed for Read and Write operations. Each drive belonging to the RAID drive is shown with it's by-id device name. The upper table represents reads, the lower table represents writes. The numbers at the top of the table columns are times in milliseconds. For example, the first column indicates 0-15 milliseconds, the second indicates 16-31 milliseconds, and so forth. The numbers below are quantities of sectors. The numbers reflected in the tables are either since the system was booted, or since the last time the tables were reset. In the example above, the first drive has a 1 in the 0-15 column in the Read table. This indicates that it has read 1 sector, and that it took between 0 and 15 milliseconds to read that sector.

Below the two tables, is a [Reset Performance Response Counters](#) button, which is used to reset the tables, and a [Return to NumaRAID GUI Main Page](#) link which returns to the Main GUI Screen. It is ideal, that before you run the test, that you left-click on the [Reset Performance Response Counters](#) button at the bottom, eliminating any accumulated numbers from previous tests or normal array operations.



CAUTION: The RAID Surface Scan is a very destructive tool.

RAID Surface Scan

	RAID Name	Type	Size(G)	Offset
Sequential Scan	LittleFootOne	Read Scan	1G	0%



CAUTION: Do not click on the Sequential Scan button yet without reading the following information .

The Raid Name [LittleFootOne] indicates which RAID is going to be tested - the drives listed in the tables. 'Type' allow you to select the test type - a Read or a Write scan.



CAUTION : A write scan will wipe out any data on the RAID being tested.

'Size' selects the amount of the RAID that will be tested - in steps of 1GB, 10GB, 100GB, or the entire RAID. Offset will let you specify a starting percent. For example, specifying 10% will mean that you want to run the test at 10% into the diameter from the outside of the drives.

Real Time Response Characteristics - In Milliseconds - Sequential Read

Drive	0-15	16-31	32-47	48-63	64-79	80-95	96-111	112-127
scsi-35000c50006969a43	11276	14872	11955	10251	2449	885	368	506
scsi-35000c50007bda8f3	22510	14243	11024	6991	596	283	110	175
scsi-35000c50006965433	17762	12034	11884	9283	1572	675	294	346
scsi-35000c50007bd5297	22598	14619	11150	6751	458	263	92	145

In this case, the numbers are low because this is a very slow array – the drives are connected to a PCI/X SAS card. In this test, using the first drive as an example, 11276 sectors fell into the 0-15ms transfer time range, 14872 fell into the 16-31ms range, 11955 fell into the 32-47ms range, and so forth.

Now as the offset changes, or if the drives are tested for larger ranges, the drives will slow down, as the heads near the inside diameter - the slowest parts of the disks. The numbers will appear to "creep right" - i.e. the left columns will start to decrease and the average will move further to the right. If you start to see a large pile of numbers in the 112-127 column, there may be a problem. In fact, if you ran a read scan across the entire RAID, and one disk had numbers only in the 112-127 column - that would be a really serious problem - go to Slots, and check the SMART for that drive to see if it is sensing anything wrong with itself - it could be near failure.

3.1.5 LUN Details

LUN Details

	Lun Name	RAID Name	Size(G)	Offset(G)
Delete	MyLun	LittleFootOne	546	0

LUN Surface Scan

	LUN Name	Type	Size	Offset
Sequential Scan	MyLun	Read Scan	1G	0%

The LUN Details screen allows you to manage the LUN, as well as run a surface scan on a single LUN. These are very powerful features currently not found on other arrays.

The top table shows LUN Details, similar to how LUN Details is shown on the RAID screen, however there is a Delete function. If you want to delete the LUN (Note all initiators and targets must be removed first), left-click on the Delete button. In this table, we see the the name of the LUN (MyLun), the name of the RAID that it is part of, the size of the LUN (in Gigabytes), and offset (Also in Gigabytes).

Below the LUN details table, is a table where you can run a surface scan of a LUN. There is no separate screen for this – the results are shown on the RAID surface scan screen.



CAUTION: A write scan will erase data on the LUN.

The controls and reports are the same as the RAID Surface scan - see the previous section for instructions.

3.1.6 CONFIG Details

CONFIG Details

Configurations

Configuration Name	
Save Current Config As	
Reload Configuration	Current ▾
Record Current Configuration to All Drives	
Recover Configuration from Drive to File	Slot0 ▾

Trace File Entries=110 Overflow=0 Total Size=481000

Trace File Actions	Type	Number of Entries
Display Trace	Commands ▾	Last 25 ▾
Capture Trace to TraceFile	Last ▾	25
Control Trace	Stop ▾	

Log File

Log File Actions
Display
Reset

[Return to NumaRAID GUI Main Page](#)

This screen is used to perform a number of utility functions. The top table of functions refers to the configuration metadata itself. The configuration information contains every piece of information about the array: RAID information, LUN information, port information, file information, sensor information, slot information, license information, parameter information, and drive information. It is stored in two places: In a file on the boot drive of the array, and also on the data drives themselves.

For added security, you can use the first function [Save Current Config As](#) to make your own backup of the configuration. Simply left-click in the text area to the right of [Save Current Config As](#), enter the file/pathname that you would like to save to, then left-click on the [Save Current Config As](#) button.

The next item in the Configurations table is [Reload Configuration](#) - this is used to either reload the "regular/current" configuration into RAM, or to load one that you saved previously. Simply select the configuration that you want to load/reload with the drop-down, then left-click on the [Reload Configuration](#) button. This is also used if you want to reload a configuration that was recovered from a drive - the configuration file recovered will show up as recslot{slot #}.xml.



CAUTION: Note that reloading the configuration unloads and reloads all of the drivers associated with the Aurora LS RAID – this will disconnect all clients!

As mentioned earlier, the configuration is also written to the data drives - if you manually want to update the configuration information recorded on the drives, simply left-click on the [Record Current Configuration to All Drives](#) button. This only takes a few seconds.

The last item in the Configurations table is the option to recover the configuration information from a single drive to a file. Select the slot number of the drive that you would like to recover the configuration from with the drop-down on the right, then left-click on the [Recover Configuration from Drive to File](#) button. The file will be saved as recslot{slot #}.xml.

The second table has to do with a Trace File:

Trace File Entries=110 Overflow=0 Total Size=481000

Trace File Actions	Type	Number of Entries
Display Trace	Commands ▾	Last 25 ▾
Capture Trace to TraceFile	Last ▾	25
Control Trace	Stop ▾	

A trace file contains internal diagnostic information, which can be used by the programmers for troubleshooting. Above the table is some information about the last/current trace. It shows the number of entries in the trace file – in the example above, there are 110 records/entries. To the right of this, it shows overflow. This indicates how many entries could not be recorded because the file became too large. On the right is the current size of the trace file (in bytes). In this example, it is 481000 bytes.

The 'Display Trace' function goes to a different screen (covered in the next section) for displaying information from the trace. You have four options here – there are two options under type (Commands and all), and two options under Number of entries (First 25 and Last 25). You can select only one option from each column, then left-click on Display Trace to go to the screen to show the results of what you selected. Under type, Commands, displays only commands, all displays commands and all other information recorded. For the number of entries. Last 25 shows the information starting with the last 25 entries of the trace file. First 25 shows the information starting with the first 25 entries of the trace file.

The 'Capture Trace to TraceFile' records the data to a file. This is usually done to retain the information from a trace prior to resetting/restarting a new one. The type function works in conjunction with the number of entries function, creating something similar to the Number of Entries function under display trace, but more flexible. You can specify "All" for type, then the number in the Number of entries field is not used – this specifies dumping all of the trace file to the data file. Otherwise you can specify First or Last, followed by the number in the next field, indicating to dump that number of entries from the start or the last of that number of entries perspective. For example, specifying First under type, then 30 under Number of entries will dump the first 30 entries to the file. Once you have made the settings that you want, left-click on the Capture Trace to TraceFile button to capture the trace to a file.

The 'Control Trace' function controls the trace. The options which appear under type change, depending on whether or not a trace is running. There's three options: Start, Stop, or Reset. Stop only appears if a trace is running, and is used to stop a trace. Start only appears if a trace is not running, and is used to start a trace. Reset only appears if a trace is running, and stops then restarts the trace in a single operation. To perform the desired action, select the action under type, then left-click on the Control Trace button.

Below the Trace table is a Log File table as follows:

Log File
Log File Actions
Display
Reset

This is used to display or reset the NumaRAID log file. Resetting the log clears the log. Display shows it. Here is a sample of what that might look like:

FILE DISPLAY

Command Results
nrcli file display errlog
Mon Sep 28 12:00:08 2009 CLI file display errlog
ok
Log
Mon Sep 28 11:59:42 2009 CLI file display errlog
Mon Sep 28 12:00:00 2009 Config Monitor Running
Mon Sep 28 12:00:08 2009 CLI file display errlog
bye
Command Succeeded
Return to Previous Page

To return to the Main GUI screen, by clicking the [Return to NumaRAID Main GUI Page](#) link at the bottom of the Config Details screen.

3.1.7 TRACE Details

TRACE Details

Trace Results 'commands' Total Entries=2048 Display=25 Current entry=

Time	Entry	uGap	uSecs	User	Lun#	Lun	CDB	Logical LBA	Length	Dirty	Status
11:59	2023	0	372	localhost	0	MyLun	READ10	2071552	1024	0	0
11:59	2024	588	5590	localhost	0	MyLun	READ10	2072576	1024	0	0
11:59	2025	50	5823	localhost	0	MyLun	READ10	2073600	1024	0	0
11:59	2026	5966	288	localhost	0	MyLun	READ10	2074624	1024	0	0
11:59	2027	346	263	localhost	0	MyLun	READ10	2075648	1024	0	0
11:59	2028	405	267	localhost	0	MyLun	READ10	2076672	1024	0	0
11:59	2029	324	271	localhost	0	MyLun	READ10	2077696	1024	0	0
11:59	2030	386	286	localhost	0	MyLun	READ10	2078720	1024	0	0
11:59	2031	345	287	localhost	0	MyLun	READ10	2079744	1024	0	0
11:59	2032	420	5184	localhost	0	MyLun	READ10	2080768	1024	0	0
11:59	2033	45	4870	localhost	0	MyLun	READ10	2081792	1024	0	0
11:59	2034	5302	287	localhost	0	MyLun	READ10	2082816	1024	0	0
11:59	2035	21	566	localhost	0	MyLun	READ10	2083840	1024	0	0
11:59	2036	701	431	localhost	0	MyLun	READ10	2084864	1024	0	0
11:59	2037	47	403	localhost	0	MyLun	READ10	2085888	1024	0	0
11:59	2038	541	279	localhost	0	MyLun	READ10	2086912	1024	0	0
11:59	2039	336	355	localhost	0	MyLun	READ10	2087936	1024	0	0
11:59	2040	494	5516	localhost	0	MyLun	READ10	2088960	1024	0	0
11:59	2041	49	5713	localhost	0	MyLun	READ10	2089984	1024	0	0
11:59	2042	5858	272	localhost	0	MyLun	READ10	2091008	1024	0	0
11:59	2043	329	272	localhost	0	MyLun	READ10	2092032	1024	0	0
11:59	2044	411	268	localhost	0	MyLun	READ10	2093056	1024	0	0
11:59	2045	324	260	localhost	0	MyLun	READ10	2094080	1024	0	0
11:59	2046	405	490	localhost	0	MyLun	READ10	2095104	1024	0	0
11:59	2047	45	361	localhost	0	MyLun	READ10	2096128	1024	0	0

Previous 25 Trace Entries

Goto Entry 2048

Switch to All display

Switch to Chart display

[Return to NumaRAID GUI Main Page](#)

The details of a 'Trace' command is very helpful to support the Aurora LS. In the example above, "Commands" and "last 25" were chosen from the Config Details screen, then a 'Display Trace' was taken to capture that data. The trace shows the last 25 low-level commands that were executed. Above the table is a description of what the trace has captured – i.e. commands or all. It shows the total number of entries, how many it is displaying, and the offset. In the table, on the left, we see the time in hours/minutes. These will almost never change from one row to the next, unless the array is idle for a long period of time, has done very few commands, or the commands are taking unusually long to execute. The entry column shows the number for the particular entry in the Trace file. uGap is the number of microseconds between commands. uSecs is the amount of time in Microseconds, that it took to execute the command. User is the originator of the command. localhost indicates that the array itself requested the command. Lun# is the logical LUN number of the LUN that the command was performed on. Lun is the name of the LUN that the command was performed on. CDB describes what command was issued, along with the length of the CDB (Command Data Block). In the first line, for example, it says "READ10" – This means the command was a read command, and the command data block for that command was 10 bytes long. To the right of this a logical LBA. This is the logical block or sector that the command was told to act on (in this case, read from). The next column is Length – this is the length of the data that the command was told to act on – in this case, it was told to read 1024 bytes. Dirty is the number of dirty segments in the cache. Status is the result of the command as reported by the device – 0 indicates that the command was successful. A non-zero number indicates the command failed.

In this case, prior to getting to this screen, we specified that we wanted the last 25 commands, and that commands were shown. If non-commands (All) was chosen, non-commands would also be in the table. If there are entries before the screen we are looking at, a button at the bottom will appear allowing you to see the “Previous 25 Trace Entries.” If there are entries after the ones shown, you will see a button allowing you to see the “Next 25 Trace Entries.” If you are somewhere in the middle, you will see both buttons, and if there are less than 25 entries, you will not see either button. Below these is a button which allows you to go to a specific entry. When you do, it will show the list of 25 entries (if there are 25), starting at the entry that you specified.

Below the Goto Entry button, is a button where you can toggle between the view of the commands, and view of all. Simply click this button to toggle between the two.

The bottom button switches to a chart display, which is explained below.

The [Return to NumaRAID GUI Main Page](#) link at the bottom returns to the NumaRAID GUI Main screen.

The Chart Display, and example shown below, shows a series of charts, graphing the information shown in the Trace Details screen. For each chart, the horizontal axis is the entry number. There are 10 charts in total. Note that the charts are showing 200 entries at any given time, as opposed to 25 entries.

The top left chart shows the logical block address (LBA) or logical position number/sector number within the RAID that the “virtual” head is positioned. In the example, it is a straight line going up to the right, because it is the tail end of a sequential read. The vertical axis is the LBA address.

The top right chart shows the transfer lengths. In my example, all of the lengths are 1024 bytes. The vertical axis is the transfer length.

The left chart in the second row indicates the access times to the cache in microseconds. The vertical axis is the time.

The right chart in the second row indicates the time it took to execute the command in microseconds. The vertical axis is the time.

The left chart in the third row shows data transfer rates. The vertical axis is in megabytes per second.

The right chart in the third row shows the command transfer rates. The vertical axis is in megabytes per second.

The left chart in the fourth row shows the write back cache usage. The vertical axis is number of write backs.

The right chart in the fourth row shows read ahead cache usage. The vertical axis is the number of read-aheads.

The left chart in the bottom row shows non-real-time commands. The vertical axis is the number of commands.

The right chart in the bottom row shows write cache saturation. The vertical axis is the number of dirty cache segments.


At the bottom of the graphs, similar to the data display, are two buttons: One allows you to go to the previous 200 entries (if there are any). One allows you to go to the next 200 entries (if there are any). Finally, there is a box you can type a number in, along with a GoTo button, which allows you to display 200 entries starting with the entry number specified.

Below these is a button which allows you to switch back to the data/text display.

To return to the NumaRAID GUI Main page, left-click on the link at the bottom, which reads [Return to NumaRAID GUI Main Page](#).

TRACE Details

Trace Results 'all' Total Entries=2048 Display=25 Current entry=2023



3.1.8 USER Details

The User Details screen is used to give the user a name, as well as assign whether or not they are real-time users. In the example above, there are two clients connected. One is connected via Infiniband, the other is connected via Fibre Channel. Also in the above example, no users or real-time users exist. Starting at the bottom of the screen, we have a table showing the status of what are called sessions. A session means a communication link has been established between the array and the client system, and that the client system is visible to the array as a potential user. In this table, Fibre Channel clients are identified under the driver column as “NumaRAID Target Driver for Atto Celerity.” Infiniband clients are identified under this column as “ib_srpt.” These drivers are the drivers which are on the array which are being used to identify the client with. The next column to the right is labeled “Target” – it pertains only to Fibre Channel clients. The target will indicate “NULL” for Infiniband clients. On Fibre Channel clients, it will indicate the physical port number on the Fibre Channel card within the array, that the client is connected to by showing “ATTOTarget{port#}.” In the example above, it shows ATTOTarget0, indicating port 0, which is the first port. The right column shows the WWN# (World-Wide Network Number). Normally, the initiator (client) is always referred to by this WWN#, but look at them – they’re long and probably impossible to memorize. The main purpose of this screen, is to assign a name that the administrator of the array can remember, to that WWN#. To assign a name to the particular item, type a name under User Name, and left-click on the Create button. Also in this table is a line for a localhost. This gives the ability for you to name the array, and mount the array on itself, if necessary.

Here is an example, using “MacPro” for the Fibre Channel user, and “Gamer” for the Infiniband user:

USER Details

Users		
	User Name	WWN
Delete	Gamer	0x0002c9030000a50c0002c90200218058
Delete	MacPro	2100001086025a20
Create		

Real Time Users	
	WWN
No Real Time Users Defined	
Create	Gamer

Session status				
	User Name	Driver	Target	WWN
	MacPro	NumaRAID Target Driver for Atto Celerity	ATTOTarget0	2100001086025a20
	Gamer	ib_srpt	NULL	0x0002c9030000a50c0002c90200218058
Create		LocalNR	LocalNRtarget	localhost

[Return to NumaRAID GUI Main Page](#)

Many things on this screen changed - Gamer and MacPro are now listed in the top table, with their names and WWN#s. You can delete either of these users by clicking Delete to the left of the name that you would like to delete.

Note that deleting the user only deletes the name given to the WWN# - nothing more.

You now also get the option in the top table, to manually enter a user/WWN#. At the bottom, under session status, we now see the users listed by name, rather than empty text boxes.

Now examine the Real Time User table. It previously showed only "No Real Time Users Defined", but now also shows a Create button, defaulting to one of the users (in this example, Gamer). Real-time users are users who get the priority over the user of the storage that they request, while the other users get whatever is left. This only matters if there is more than one

user. To make a user a real-time user, select their name from the drop-down on the right, then click the Create button on the left.

Using “Gamer” as an example, the middle table will change as follows:

Real Time Users		
	Name	WWN
Delete	Gamer	0x0002c9030000a50c0002c90200218058
Create	MacPro ▾	

If you wish to make the user not-real-time, left-click on the Delete button to the left of it's name. Note that there can be multiple real-time users – for example, I could also add “MacPro” as a real-time user.

To return to the Main GUI screen, left-click on the Return to NumaRAID GUI Main Page link at the bottom of the User Details screen.

3.1.9 PARAM Details

PARAM Details

Streaming Read Parameters

	Parameter Name	Parameter Value
Update	Maximum Read Ahead Distance in 128k Stripes	24
Update	Stripes Required in Memory before Read Ahead Allowed	24
Update	Maximum Read Ahead Commands Outstanding	8
Update	Number Of Stripes in Each Read Ahead Request	8
Update	Enable Random Reads	Yes

Streaming Write Parameters

	Parameter Name	Parameter Value
Update	Cache Flush Percentage Threshold (0-100)	10
Update	Maximum Write Back Requests Outstanding	8
Update	Number Of Stripes in Each Write Back Request	8

Real Time Parameters

	Parameter Name	Parameter Value
Update	Percent of Cache Available To Non-Real-Time Writes	50
Update	Percent of Cache Available To Real-Time Writes	75
Update	Max Data Rate of Non-Real-Time Requests(MB/Sec) 0 for no limit	0
Update	Max Number of Non-Real-Time Requests 0 for no limit	4
Update	Reconstruct In Advance of Drive Completion	No
Update	Reconstruction Priority (from 0 to 100)	0

Silent Data Corruption Detection and Correction

	Parameter Name	Parameter Value
Update	Enable PQ Verification	No

Logging Level

	Parameter Name	Parameter Value
Update	Internal Diagnostic Message Level	Disabled
Display	Internal Diagnostic Message Log	

[Return to NumaRAID GUI Main Page](#)

The PARAM functions are for setting or viewing global array parameters. Each row in each table except for the last row of the last table, shows a parameter and value. Should you need to change this value, you would alter the value on the right, then click the corresponding update button on the left. It works the same way for changing every parameter. So here are the parameters and what they mean/do:

Maximum Read Ahead Distance in 128k Stripes: When you playback video for example, you are essentially doing one large sequential read. To make playback smoother, the array can be set to read more of the file than the position that the client computer is currently requesting. This is called a read-ahead cache. The cache is only selectable in 128KB increments, and the value here is the number of 128KB blocks to use (The blocks are referred to as stripes, because they go across all of the drives in the RAID). The default value is 24. This allows the computer to read 3MB ahead. So, for example, if you were playing a standard-definition video file, which plays relatively slowly in relation of the array, when the computer playing the video starts playing at 12MB of the file (for example), the array has already read the next 3MB, and is ready to play up to 15MB, without doing any disk activity. As the computer plays through this cache, it is refreshed with new data as necessary. Making this setting it too high would cause kind of a stopping/starting of data reading on the array, and setting it too low would render the cache not as effective.

Stripes Required in Memory before Read Ahead Allowed: This is the amount of sequential data that must be read in order to trigger the read-ahead cache above. The default value, 24, (using the same stripe value as above 128KB), means that the client must request 3MB of sequential data in order to activate the cache. Setting this value too low would force the array to re-cache over and over as fragmented files occur. Setting it too high might force it not to cache something that otherwise would benefit the client.

Maximum Read Ahead Commands Outstanding: While the array will appear to be sending and receiving data, the client is also sending commands to the array to tell it to read or write data. The client, for example, might send a request to the array to send back (read) 1MB of data, however before the array has finished, the client might send a request to the array to send back another 1MB. This is happening anywhere up to millions of times per second. This setting controls how many of those commands will be buffered at a time. The default value of 8 is good for most cases. Setting the number too low may result in jerky playback - i.e. the computer sends a request, the array sends back the data, then waits for the next request. Setting it too high would just waste memory.

Number of Stripes in Each Read Ahead Request: This can control the size of each request. The default value is 8 (x 128KB) which is 1MB. This keeps the data coming from the array at a consistent rate - i.e. if the requests from the client were not limited, the requests might be uneven, possibly interrupting playback for other clients.

Enable Random Reads: The array is capable of applying the read-ahead cache to non-sequential sectors/stripes. The default value enables this. If it is disabled, the read-ahead will only apply to sequential reads where the sectors/stripes themselves are sequential.

Cache Flush Percentage Threshold (0-100): This controls how often when writing, that the cache should write its contents to disk and empty itself. The default value is 10 (%), which means that when the cache is at least 10% full, it should empty. The cache size which was chosen when the RAID was created has a direct bearing on this setting. For example, if you used a cache size of 3GB, and this value is set to 10, then the write cache will flush when it is roughly 300MB full. The default number is fine in most cases. If you set the number too low, you will disable the effectiveness of the write cache, as it will be emptying more often. If you make it too high, you risk having to wait for a larger cache flush.

Maximum Write Back Requests Outstanding: Just as you can control how many commands the read will buffer, you can also control the amount of commands that the write will buffer. The default value of 8 is good for most cases. Setting the value too low or too high may result in dropped frames on capture because either you are not allowing the client computer to send enough write commands, or are accepting too many. Setting the value too high will waste RAM.

Number of Stripes in Each Write Back Request: This setting controls a limit on the amount of cache to use for each write command from a client. The default value is 8, which is 1MB. This is fine in most cases. Making the value too low would limit the cache too much. Making it too high would probably just waste RAM.

Percent of Cache Available to Non-Real-Time Writes: This applies to the real-time users. You can actually dial-down the cache for writing for non-real-time users. This value is a percentage. The default value of 50, indicates that real-time users only get a maximum of 50% of the cache. Setting this value too high would render this setting useless. Setting it lower would further limit the cache for non-real-time users. Keep in mind, this setting only applies to non-real-time users - see below for real-time users. Note that this setting applies globally to all non-real-time users.

Percent of Cache Available to Real-Time Writes: This is the same as above, but only applies to real-time users. The default value of 75 indicates that a real-time user gets 75% of the cache for writes. Setting the value higher could impact non-real-time users more. Setting it

lower gives up some of the cache to the non-real-time users. It's almost the opposite of above. Note that this setting applies globally to all real-time users.

Max Data Rate of Non-Real-Time Requests (MB/SEC) 0 for no limit: This allows you to limit the bandwidth of non-real-time users in megabytes per second. It is used to free up bandwidth for real-time users as well. The value entered here is in megabytes per second. The default value, 0, does not limit the maximum data rate for non-real-time users.

Max Number of Non-Real-Time Requests: Another way of limiting non-real-time users is to limit the amount of read/write commands they can send. Note that this setting affects all non-real-time users. The default value is 4. Setting the value lower would further limit non-real-time users. Setting it higher would cache more requests.

Reconstruct in Advance of Drive Completion: If a drive isn't performing as well as the rest, this option is used to base the data on the parity, instead of the data returned from the drive. In many cases, this can compensate for a slow drive. This option is disabled by default.

Reconstruction Priority (from 0 to 100): The array is capable of reconstructing while it is being used. This value controls the balance of priority given to reconstruction versus the data access. The default value is 0, which means reconstruction is only performed when the array is idle. If you set it to 100 (which is definitely not recommended), the array would run very slowly to the clients, while reconstructing at full speed. So as an example, consider a value of 10 - This would mean that the array would spend 10% of it's time while being accessed, doing reconstruction. The value is up to you - the more time and/or speed you can sacrifice while the array is being used to reconstruction, the faster the reconstruct will complete.


Enable PQ Verification: Default is No. This value is a form of error-detection and correction (Raid-6 only). If this value is enabled, while the array is reading, it will compare the data read against the two parity generators - there has to be a 3-way match between the data and each of the two parity generators. If there isn't, the data from the parity generators is used instead of the data from the drive in question - this substitution is made in real-time. So basically, if the array detects something wrong in the data, it corrects it. Enabling this option might affect read transfer rates.

Internal Diagnostic Message Level: More explicitly, this value determines what you want the internal diagnostics to log. Here are the values and what they do:

Disabled	Do not log anything.
Requests	Only log read/write requests.
State Started	Only log state engine starts.
State Ended	Only log state engine completions.
BIO Started	Only log Block I/O starts.
BIO Ended	Only log Block I/O completions.
Cache	Monitor the cache.
Debug	Monitor debugging.
Performance	Monitor performance.
Target	Monitor targets.
Silent Data Corruption	Monitor for the problem described under PQ Verification.

The [Display](#) button on the bottom of the table, displays the diagnostic message log. Note that it is important that this function only be used when directed to do so, and it must be disabled when not in use, otherwise it would fill up the boot drive. Here is a sample of what part of that log/output would look like:

PARAM



Command Results
dmesg

```
ok
session 0xffff810121a12000
NumaRAID: In scst_set_cmd_error, line 2687: CPU 1 scst_set_cmd_error(ffff810213760010, 5, 37, 0)
NumaRAID: In scst_set_cmd_error_status, line 2659: CPU 1 scst_set_cmd_error_status(ffff810213760010, 2)
NumaRAID: In scst_rx_cmd, line 1070: CPU 1 No LUN 133 for session 0xffff810121a12000
NumaRAID: In scst_set_cmd_error, line 2687: CPU 1 scst_set_cmd_error(ffff810213764010, 5, 37, 0)
NumaRAID: In scst_set_cmd_error_status, line 2659: CPU 1 scst_set_cmd_error_status(ffff810213764010, 2)
```

The [Return to NumaRAID GUI Main Page](#) link at the of the Parameters Details screen will return you to the NumaRAID Main GUI Screen.

3.1.10 DATARATE Details

For ease of discussion DATARATE details functions and options will be discussed by section.

DATARATE Details

DataRate Options

	Sample Source	Name
Chart	NumaRAID	Device
Chart	User	ATTOTarget0:MacPro ▾
Chart	RAID	SmallFoot ▾
Chart	LUN	MyLun ▾
Chart	Target	ATTOTarget1 ▾

At the top of this screen is a series of options for controlling what charts you see at the bottom. You would select what you would like to view on the right, then click the corresponding Chart button on the left to see the charts below for that selection. The options are as follows:

NumaRAID – Device: This shows graphs pertaining to the entire Aurora RAID Array.

User: Allows you to see graphs pertaining to I/O for a particular user.

RAID: Allows you to see graphs pertaining to a particular RAID.

LUN: Allows you to see graphs pertaining to a specific LUN.


Target: Allows you to see graphs pertaining to a specific Fibre Channel Target/port.

For these examples, the default (NumaRAID – Device) is used.

There are 6 sets of graphs in each set. The upper right of each shows information pertaining to the current minute. The upper left shows the previous minute. The middle right shows the current hour, the middle left shows the last hour. The lower right shows the current day, and lower left shows the last day. On each set of charts, read information is in green color, and write information is in red.

Legend - Reads in GREEN Writes in RED

Data Rate for Object 'NumaRAID' Name 'Device'



The first group of charts is for data rates. Vertically, the rate is shown in megabytes per second. If you examine the example, the array spent approximately 57 seconds of the last minute, doing a data rate test which yielded a result of about 410 megabytes/second. This test proceeded through the next 55 seconds or so into the current minute. If you look at the middle-right chart, you can see that the test took roughly 2 minutes to perform.

The second set of 6 charts shows response times:


Response Time



In this set of graphs, going across, we don't see the actual time as in the first set of graphs, but divisions of times. Vertically, it is showing the number of commands executed. Horizontally, it is how long each command took during that time period. So for example, in the upper right chart, we see four bars: The left bar shows that there were about 39,000 commands executed which took 100 microseconds to execute. The middle bar shows that there were about 2,900 commands executed which took 1 microsecond to execute. The third bar shows there were about 3,000 commands executed which took 10 milliseconds to execute, and the fourth bar (almost not visible) indicates maybe several hundred commands which took 100 milliseconds to execute. In this example, this is a good array, and these are values like you would find on good arrays – big bars on the left, little or no bars on the right.

The third set of graphs shows transfer sizes:

Transfer Size



[Return to NumaRAID GUI Main Page](#)

This set of charts is showing the number of commands, versus the transfer size at the bottom. If you use only one application to access the array, what you would like to see here is a single bar, as far to the right as possible. This indicates that the array did a lot of large transfers, which were all equal in size. Going vertically is the number of commands/transfers, and horizontally is the transfer size. In my example, there were about 96000 transfers performed, each of which was 512 kilobytes in size.

If you would like to return to the NumaRAID Main GUI Screen, left-click on the [Return to NumaRAID GUI Main Page](#) link at the bottom of the Data Rate Statistics Screen.

3.1.11 SLOT Details

SLOT Details

	Number	Vendor	Model	Rev	Capacity(G)	Device	Linux Device	Status
Smart	0	SEAGATE	ST3146855SS	0002	146	scsi-35000c5000696adf3	sdb	ok
Smart	1	SEAGATE	ST3146855SS	0002	146	scsi-35000c50007bd20ef	sde	ok
Smart	2	SEAGATE	ST3146855SS	0002	146	scsi-35000c5000696997f	sdh	ok
Smart	3	SEAGATE	ST3146855SS	0002	146	scsi-35000c5000696991f	sdh	ok
Smart	4	SEAGATE	ST3146855SS	0002	146	scsi-35000c50002f2e48f	sdh	ok
Smart	5	SEAGATE	ST3146855SS	0002	146	scsi-35000c50006965a83	sdf	ok
Smart	6	SEAGATE	ST3146855SS	0002	146	scsi-35000c5000696ae97	sdi	ok
Smart	7	SEAGATE	ST3146855SS	0002	146	scsi-35000c500069698fb	sdl	ok
Smart	8	SEAGATE	ST3146855SS	0002	146	scsi-35000c50006969a43	sdd	ok
Smart	9	SEAGATE	ST3146855SS	0002	146	scsi-35000c50007bda8f3	sdg	ok
Smart	10	SEAGATE	ST3146855SS	0002	146	scsi-35000c50006965433	sdj	ok
Smart	11	SEAGATE	ST3146855SS	0002	146	scsi-35000c50007bd5297	sdm	ok
Smart	12	SEAGATE	ST3146855SS	0002	146	scsi-35000c50007bceb93	sdn	ok
Smart	13	SEAGATE	ST3146855SS	0002	146	scsi-35000c50006965a33	sdo	ok
Smart	14	SEAGATE	ST3146855SS	0002	146	scsi-35000c50006965a77	sdp	ok
Smart	15	SEAGATE	ST3146855SS	0002	146	scsi-35000c50007bd30a3	sdq	ok

[Return to NumaRAID GUI Main Page](#)

The slots are the physical drive bays (or drive slots) located in the array itself. It's important to note that the slot number does not necessarily correspond to the logical position of a drive within a RAID. For example, you could have a chassis with 24 slots, but have (2) 12-drive RAID's defined, each of which, with a drive 0, 1, 2, etc., but there would only be one slot 1.

For each slot in the array, we see the slot number, drive manufacturer, model number, firmware revision, capacity (in Gigabytes), the by-id device name, Linux short device name, and current status of that slot. The [SMART](#) button to the left of each drive takes you to the SMART details for that particular drive below. When you are finished and wish to return to the NumaRAID Main GUI screen, you can left-click on the [Return to NumaRAID GUI Main Page](#) link at the bottom of the Slot Details screen.

Modern hard drives have sensors within them that can log and detect problems, which can cause a drive to prematurely fail. They also run self-diagnostics and record the results. The output of SMART is different for a SATA drive versus a SAS drive. Here are some of the things that SMART might show:

Device: Shows the manufacturer, model number, and firmware revision for the device.

Serial Number: Is the serial number: Note that the actual serial number is just the rightmost 8 characters. The rest of the string is a manufacturer-unique ID.

Device Type: Shows the type of the device.

Transport protocol: Connection type - i.e. SAS or SATA

Local Time: Shows the time that this command was executed.

SMART Feature: Indicates whether or not the drive supports the SMART feature, and whether or not it is enabled.

Temperature Warning: Indicates whether or not a temperature warning is enabled or disabled.

Overall Health: Indicates the drive Health at the time this command was executed.

Current Drive Temperature: This is the temperature (in Celcius) at the time the command was executed.

Drive Trip Temperature: Indicates the maximum internal temperature that the drive ever recorded.

Elements in Grown Defect List: The drive keeps track of different areas that it can not write to. These are called "surface defects." There are two defect lists: One is the Manufacturing Defect List, which contains defects that were found when the manufacturer tested the drives. This list is fixed and never changes. The other list is called a grown defect list, which is a list of defects that occurs after the drive leaves the manufacturer. This list only gets bigger, hence the "grown" name.

Vendor Cache Information: This is just a category heading which describes the next 5 lines.

Blocks Sent to the Initiator: In the case of SAS, the host adapter channel is called an initiator, while the drive itself is the target. This line indicates the number of blocks of data sent to the initiator – in this case, the blocks are 512 bytes (sectors), however they may or may not be data from the disk – they could also be SMART data such as the one which was requested here. Most of the time, these are drive data sectors, so in general, this is the number of sectors that has ever been read from the drive.

Blocks Received from the Initiator: In general, this is the number of sectors written to the drive.

Blocks Read from Cache and sent to the Initiator: This is an indicator of how efficient the caching is on the drive. If the computer (initiator) requested the same block twice, and it happened to be in the cache of the drive, then the drive would not have to read it again from the disks, so in general, this number would be the same or always higher than the Blocks sent to the Initiator. The higher the number goes, it means the less work the heads on the disks have to do.

Number of Read or Write Commands who's size <= Segment Size: The drive only sends data to the computer in groups of blocks, into an area of the cache, called a cache segment. If the commands being sent or the data being sent back is smaller or the same size as a cache segment, it would register here. This number doesn't necessarily indicate something good or bad – just a number of commands sent which were not the same size or smaller than the cache segment – most are not.

Number of Read or Write Commands who's size > Segment Size: This indicates data or commands which had to be broken up into multiple transfers to send to the drive or the computer. This doesn't mean anything good or bad.

Vendor (Factory) Information: This is a category heading for the next two lines.

Number of Hours Powered Up: This indicates how long a drive has been powered up (in hours), regardless of whether or not it was reading or writing – even just sitting idle counts as being powered up. In fact, if the drive had power and was put to sleep, it would also be counted here.

Number of Minutes until next SMART test: The drive has two diagnostic tests. One is a quick test, which only takes a few seconds, and is run by the drive itself (if not manually triggered). The other is a full surface scan, which is only initiated by the user. In this example, there is 1 minute until the drive is going to run the quick test on itself. The quick test is how the drive updates this information.

The next section shows the Error Counter log. The output, when viewed with a fixed-space font, forms a table – here is a sample of what that table might look like:

```
Error counter log:
      Errors Corrected by
      ECC          rereads/
      fast | delayed rewrites
read: 130744731    235      0    130744966    130744966    8302.908      0
write:      0      0      0      0      0    11336.165      0
verify: 5990726      0      0    5990726    5990726      0.000      0
```

Definition of log entries:

'read' row is showing numbers relating to reads.

'write' row shows numbers relating to writes. Most of the write row will always be 0, because this particular drive does what are called blind writes (i.e. Isn't capable of detecting errors on writes without a verify or read)

'verify' row shows numbers relating to verifies (which are writes followed by reads to check the data).

The first two columns are errors corrected by ECC (Error Correction and Control). With ECC extra bits are sent with the data which provide parity for the data. If the parity doesn't match the data, it is corrected by the processor on the drive. The third column shows errors which were corrected by rereads (Where the drive had to reread the sector to get the data), or rewrites (Where the drive had to write the sector more than once, based on a verify failure). The forth column shows the total numbers of errors corrected (i.e. The sum of the first three columns). The fifth column shows how many times it had to call the error correction algorithms (whether or not the errors were corrected) – kind of also like a sum of the first three columns. The sixth column indicates how many Gigabytes have passed through the error-checking algorithm. In this case, a little over 8.3TB was processed. Finally, the right column is number of errors which could not be corrected either with ECC or with rereads/rewrites.

The final two lines are: GLTSD, which records multiple test results (it should be disabled), and finally, the long (extended) self-test duration, which indicates the amount of time in seconds and minutes that it took the last time it ran the long self-test. This is a good indicator of how long futures tests would take to run. In the example, the test took about 63 minutes to run, which is very good for a 1TB SAS drive.

The following is a sample output of the SMART command from a SAS data drive:

SLOT

Command Results

```
smartctl -a /dev/sdb
```

ok

Device: SEAGATE ST3146855SS Version: 0002
Serial number: 3LN2BD3T000098124404
Device type: disk
Transport protocol: SAS
Local Time is: Tue Sep 29 11:03:53 2009 CDT
Device supports SMART and is Enabled
Temperature Warning Disabled or Not Supported
SMART Health Status: OK

Current Drive Temperature: 38 C
Drive Trip Temperature: 68 C
Elements in grown defect list: 0
Vendor (Seagate) cache information
Blocks sent to initiator = 797365557
Blocks received from initiator = 1301127888
Blocks read from cache and sent to initiator = 2557826457
Number of read and write commands whose size <= segment size = 184154329
Number of read and write commands whose size > segment size = 41817
Vendor (Seagate/Hitachi) factory information
number of hours powered up = 2036.85
number of minutes until next internal SMART test = 13

Error counter log:
Errors Corrected by Total Correction Gigabytes Total
ECC rereads/ errors algorithm processed uncorrected
fast | delayed rewrites corrected invocations [10^9 bytes] errors
read: 35983313 0 0 35983313 35983313 32337.897 0
write: 0 0 0 0 0 1746412159.761 0

Non-medium error count: 2

[GLTSD (Global Logging Target Save Disable) set. Enable Save with '-S on']
No self-tests have been logged
Long (extended) Self Test duration: 1367 seconds [22.8 minutes]

Command Succeeded

[Return to Previous Page](#)

3.1.12 SENSOR Details

SENSOR Details

Name	Value	Status	Lower Limit	Upper Limit
3.3V	3.3	ok	2.9	3.6
12V	11.81	ok	+10.75	+13.25
5V	4.94	ok	+4.64	+5.65
5VSB	4.92	ok	+4.64	+5.65
Batt	3.18	ok	+2.99	+3.66
IntRightFan	4591	ok	712	18000
IntMiddleFan	4607	ok	712	18000
IntLeftFan	4655	ok	712	18000
EnclosureTemp	37	ok	+0	+50.0

[Return to NumaRAID GUI Main Page](#)

A sensor is usually a chip, optical sensor, switch, or specialized resistor located inside the array, used to detect the status of a component, such as detecting a voltage, fan speed, or temperature. This screen allows you to see the various sensors, and the range for each. A sensor which goes out of this range could indicate a component which either has failed or which may fail soon.

For each sensor, we see the sensor name, its current value, and a status indicator which indicates whether or not it is inside of the range. The lower limit and upper limit define the range. Here is an explanation of the sensors listed above:

3.3V: This is the +3.3V power output as seen from the motherboard. This voltage is especially important for the CPU.

12V: This is the +12V power output as seen from the motherboard. This voltage is especially important for powering the motors on the hard drives as well as the fans in the system.

5V: This is the +5V power output as seen from the motherboard. This voltage operates the majority of electrical circuits within the system.

5VSB: This is the +5V Standby power output as seen from the motherboard. The main use of this is it powers the circuitry necessary to turn on the system. It also powers the IPMI card (if installed).

Batt: This is the voltage of the CMOS battery. This battery retains the settings for booting the array when the system is off or unplugged.

IntRightFan/IntMiddleFan/IntLeftFan: These are the main system cooling fan speeds. On the array in the example above, there are three internal fans – they are located internally in the center of the array, one on the right, one on the middle, and one on the left. In this example, the fans spin at a maximum of about 4,600 RPM. Other systems may have more fans, and some systems have fans which spin as fast as 11,000 RPM.

EnclosureTemp: This is the temperature as measured at the motherboard – usually with a sensor located near the card slots.

Left-click on the [Return to NumaRAID Main GUI Page](#) link at the bottom of the Sensor Details screen to return to the Main NumaRAID GUI screen.

3.1.13 ADAPTER Details

ADAPTER Details

NumaRAID Network Management Ports

	Port Name	Current IP	DHCP	IP Address	Netmask
Update	eth0	192.168.1.91	y	192.168.1.91	
Update	eth1	192.168.1.91	y	192.168.1.91	

NumaRAID Fibre Channel Ports

Model	WWNN	Link Status	Link Speed
ATTO Celerity FC-84EN	2000001086022340	Up	8GB/s
ATTO Celerity FC-84EN	2000001086022341	Down	
ATTO Celerity FC-84EN	2000001086022342	Down	
ATTO Celerity FC-84EN	2000001086022343	Down	
Update Optional FC Card Parameters			

NumaRAID Infiniband Ports

Port	Physical State	State	Rate
1	4: ACTIVE	5: LinkUp	20 Gb/sec (4X DDR)
2	1: DOWN	2: Polling	10 Gb/sec (4X)

[Return to NumaRAID GUI Main Page](#)

This screen shows a lot of information. It shows Ethernet ports, Fibre Channel ports, and Infiniband Ports. (Note: In the example above, one Fibre client and one Infiniband client are shown).

In the top table, we see the Ethernet ports which can be used to remotely manage the array. The current port name and IP address are shown for each port. In the DHCP dropdown, “y” indicates that DHCP is being used. If you wish to enable DHCP, change the dropdown to y, clear the IP address and subnet mask on the right, then left-click on the Update button on the left. If you wish to set a static IP address, change the DHCP dropdown to n, type the IP address and subnet mask in the fields on the right, then left-click on the Update button on the left.

The middle table shows information relating to Fibre channel. The model of each port is shown, along with it's WWN#, the Link status, and link speed. The text field at the bottom along with Update Optional FC Card Parameters is used to change special settings on the Fibre Channel card within the array.

The bottom table shows Infiniband-related information. Going from left to right, you can see the port number, physical state, port state, and data rate.

The Return to NumaRAID GUI Main Page link at the bottom is used to return to the NumaRAID Main GUI screen.

Section 4

Troubleshooting Guide

4.0 Troubleshooting Aurora LS

This section contains typical types of common errors a list of common error messages and their meanings as well as corresponding tips on how to resolve the underlying problem. If your error message is not listed here please contact Aurora LS support and service team (see section “help” above). Our staff will help you find a solution.

Rorke Technical Support email support is available at techsupport@rorke.com or is available 9am-5pm five days a week by phone at 800 328 8147.


4.1 Chassis Status Indicators

The front of the Aurora LS has some indicators that can help determine basic problems with the unit.

Front Operator Panel



Next to the front operator panel Power and Reset switches is the Power LED. This illuminates when power is on. Next to the power LED is the Aurora LS OS DOM activity LED for the internal boot drive. This LED will light intermittently during normal operation. Next to the DOM LED are two network LEDs. These LEDs will light when there is activity from the ports they correspond to on the rear. Next to these is a temperature warning LED. If the temperature inside the system becomes too high, this LED will illuminate. Next to the temperature warning LED is a power supply / fail warning LED. If there is something wrong with the power, this LED will illuminate.



Each Aurora LS drive canister has 2 LEDs. The top LED flashes Blue and indicates the drive is functional. The bottom LED shows Red when the drive has been detected as failing to operate properly. The bad drive will cause the RAID to show a “degraded” status in the GUI and its location in the RAID will have a Red ‘FAILED’ indication.

4.2 GUI status indicators

The Aurora LS has many background sensory programs that pass data to the GUI and simplify the ability to check status and determine where problems are. Use of the RAID, DRIVE, ADAPTER and SENSOR details will give you good indications of how each major component is working.

4.3 Power System

The power system itself has several components, depending on the type of power system used. Here are power system components that Rorke has had experience with:

A single ATX-style power supply containing a single fan, non-removable power supply, with no direct status monitoring. Single power cord.

A single removable power supply system, with a fan at either end of the power supply module, and a DC power distribution board that the power supply module plugs into, with status monitoring. Dual power cord.

A dual-redundant power supply system, with a fan at either end of the each power supply module, and a DC power distribution board that the power supply module plugs into, with status monitoring. Dual power cord.

While these power system configurations may seem drastically different, there are a large number of components in them which are common to all three. The motherboard/array currently does not monitor the output of the power supply status cable – it looks directly at voltages. Here are some components, along with possible problems/fixes:

Power cord: The majority of power problems that people have are from things which are outside of the system. On any power system, if there's no power going in, it will simply not turn on. If the cable itself is damaged, it also may not turn on. If the power source is not providing power (i.e. the wall outlet), it will not turn on, and finally, if either plug on the power cable is damaged, it may not turn on. One other thing worth mentioning along these lines is electrical sparks coming out of the power connection on the power supply when it is connected – this is typically due to a worn-out power cord or damaged receptacle on the power supply. If sparks or smoke comes out of the power supply itself, it could be a problem with the power supply – unplug it immediately in either case. On a dual-power supply system, if one power supply isn't getting power for whatever reason, it will not register to the array as a power supply failure, as the power supply actually is working, but is not getting power. Of course, if neither power supply is getting power, the problem is more likely outside of the array.

4.4 Using GUI for FAN problems

The fan(s) in the power supply (or power supply modules) are temperature-controlled. The fans will operate at approximately 50% of their speed when the temperature is low, and at full speed if the temperature becomes too great. Several things can happen with the fans: If the bearings break down inside, they will stop spinning. If the blades break, they will stop spinning. If the fan motor breaks down, they will stop spinning, and if they get fouled with enough debris, they will stop spinning. If a fan starts making an unusual noise, it is a typical symptom of one of these problems. If this is the case, you do not want to ignore it. If the fan fails, a power supply failure itself, may be imminent. It can be somewhat challenging to hear the power supply fans over the noise of the main system fans – when you first plug in the power supplies with the system off, you should be able to hear the power supply fans at low-speed. In most cases, the power supply fan itself is not field-replaceable. If the power supplies are removable modules, replacing the module replaces the fan.

4.5 Using GUI for Power Supply problems

On a fixed ATX power supply, if a cable is frayed, it can be shorting something to ground. Also, it's possible for the connectors to be damaged (from repeated plugging), and aren't effective enough in contacting the motherboard. If a cable is broken, that could be a problem. Typically, the symptoms you would be looking for on a power supply are unusually low or high voltages (or both). The voltages read with by the Aurora LS's sensors are on the motherboard – if these voltages are not correct, it could also indicate a power supply problem. On a system with redundant power supplies, the power load is shared between the power supplies, so if the voltages are off, it could indicate a problem with one power supply, both, or the DC power distribution board. On systems with removable power supplies, there is usually a buzzer on the DC power distribution board which sounds if there is a voltage problem. Again, if there is no power going in to one power supply on a dual-power supply system, the buzzer may not sound, as there is no problem with the power supply – the DC distribution board is just sending out power from one power supply instead of two. Systems with removable power supplies have card-edge connectors which contact the DC power distribution board. If this card-edge connector is oxidized, scratched, or otherwise broken, it could cause a problem.

4.6 DC Power Distribution problems

On systems with removable power supplies, this is the board that the power supplies plug into. Systems with single power supplies have less-complicated DC Power Distribution Boards than ones with redundant power supplies. This is because on the ones with redundant power supplies, the board has to tolerate power surges if a power supply is hot-plugged. The board is fairly simple – it usually either works or it doesn't. The connections to the motherboard are prone to the same problems that the fixed power supplies have, but additionally, they possess a delicate communication cable which relays power supply status information to the motherboard. It is possible for the connector(s) which contact the power supplies to be broken as well – especially if someone tries to force a power supply in upside-down.

4.7 Chassis Problems

The chassis is an electromechanical system itself, which could present a myriad of problems as follows:

Air Intakes/Exhaust: These should be periodically cleaned, as their blockage could generate unnecessary heat inside the array.

Rack Mounting: On many of the chassis we've used, there are problems associated with the weight of the unit when used in a Rack configuration. The rack mounting system typically starts at the chassis itself, with a series of tangs which are punched out of the metal. In a lot of cases, these can become bent, making it difficult to attach the rails – you can bend the tangs out, but it should only be by enough to get the rail on – overbending it will cause the rail to jam when the unit is rack-mounted. The slides which attach to the sides have to go on particular sides and with a particular orientation. Currently, on chassis we've used, it isn't possible to install the slides with an incorrect orientation unless they are on the wrong sides. On the front of the chassis are a pair of rack ears. These ears are held to the chassis using screws which go into the chassis by an amount less than 1/16", and are not made to take any weight whatsoever.

MTP: The left ear on most chassis, also contains electrical connections between an MTP (Mapping/Test Panel) on the ear, which turns on or resets the power, and provides LED status information, and connects to the motherboard. On the front of the ear is a handle – most are connected with sub-standard screws which only extend into the handle by 1/8" – again, these can not take any weight. The MTP electrical connection is much more complex than it looks. Inside the rack ear is a small circuit board – on this board is a connector which is attached to a flat ribbon cable. The connector can be opened and ribbon cable removed, but it is very difficult to reassemble. The ribbon cable passes through a hole in the chassis (and can be easily damaged by metal cutting into the cable), to another circuit board inside the chassis. This inner circuit board also has a connector for the ribbon cable which can be opened/closed, then it is attached to another removable cable which goes to the MTP connector on the motherboard. The desktop chassis also contains an MTP, but is it not as delicate, easy to break, or as complex as the ones on the Rack enclosures.

Chassis Construction/Bulkheads/Air Baffles: Many of the chassis used aren't just a simple piece of metal bent into the shape of a PC. The rack-mount chassis, for example, are no less than 3 layers of metal at almost any given spot at the front, 2 at the bottom where the motherboard is, and sometimes 2 at the rear. It is possible to disassemble these layers, however the correct tools and replacement parts must be used. Most chassis have an inner bulkhead, separating the front of the chassis from the rear of the chassis, typically holding the central fans. The bulkhead is removable to allow easier access to many of the components. Finally Air Baffles: These provide directed cooling at specific components, and some provide protection for more delicate internal components. On some of the rack chassis, there is an air baffle covering the DC power distribution board. This is strictly to provide airflow while protecting the delicate components on that board. It can be removed if necessary, but should be replaced when done. Finally, there is usually a main air baffle in the system, directing air from the fans across the CPU and RAM. If the system has a Nehalem 900-series CPU, it isn't currently possible to use the air baffle, because the CPU fan required by Intel is too tall.

Mounting Hardware: While it is not likely that a piece of mounting hardware will fail in the field, one problem was discovered when developing prototypes: Not all motherboard standoff positions are used in the chassis for any given particular motherboard. If a standoff is placed in a position where there is no corresponding hole in the motherboard, it can short part of the motherboard to ground which wasn't intended, leading to possible damage or a blank screen on bootup.

Environment/Care: Environment can play a large factor in the lifespan of the array. The two harshest environments are near beaches, and in climates with high humidity. Rust forms as the result of a chemical reaction, where electrons leech out of the iron in the chassis, into the surrounding oxygen. Water and salt accelerate this reaction because they contain minute traces of electrolytes. Rust can be removed via the use of Royal Naval Jelly. But bear in mind,

if there's rust on the outside, electronic components on the inside could also be rusting – and those can't be cleaned with the Royal Jelly.

4.8 Motherboard problems

Connectors: As with the plugs which plug into them, many connectors can be damaged – especially SATA connectors on the motherboard. Here are the various connectors used and considering which could be damaged: LED/switch/Chassis connections, IPMI socket, RAM sockets, CPU sockets, PCI/PCIe slots, power connections, fan connections, SATA connections, and I2C connections (to power supply or to LEDs).

i801: The motherboards we've tested, have Intel i801 chips used for the sensors. While this is a fairly reliable chip, the symptom you might see if it fails is that all of the sensors will go dead simultaneously (Assuming there is no software problem), and/or the chip can't be found by the computer.

Northbridge: The Northbridge controls higher-speed functions of the motherboard, such as the on-board VGA (ATI ES1000 or Matrox G200) and RAM. If the on-board VGA dies, the unit is still capable of being operated remotely, however the only fix is to replace the motherboard. Note that on some motherboards, the Northbridge also controls the PCIe slots.

RAM: RAM can fail. If the amount of memory is suddenly decreased, it could indicate a problem with one or more of the memory modules. If the module is intermittent, try swapping around the modules and see if the problem goes away. If the module failed completely, the best way to troubleshoot it is to try swapping the modules one-at-a-time.

Southbridge: This chip controls the slower-speed functions of the motherboard, such as PCI/32, PCI/x, serial/parallel ports, power management, Ethernet, USB ports, and interfaces with the real-time clock. Typically, if a Southbridge dies, then entire motherboard doesn't function.

CPU: If you have a motherboard with multiple CPUs, if one CPU goes out, the system will typically lock up until it is rebooted, at which point, only one CPU might come up. See also fans, below.

Chassis/CPU/Chipset Fans: It is important to keep an eye on the chassis fans, as they not only cool the drives, but also play a part in cooling the motherboard, CPU, and RAM. There also may be, depending on the motherboard, a fan on the Northbridge or Southbridge chip, as well as a fan directly on the CPU. If a chassis fan fails, you should see it in the NumaRAID GUI, however if a chipset or CPU fan fails, a typical symptom is spontaneous rebooting of the array (Not related to software).

IPMI Card/On-Board: Typically, either the IPMI card works or it doesn't. If an IPMI card fails, it will show a host of symptoms, such as not appearing in the BIOS, or it's Ethernet port or virtual disk not showing up in the OS. However, if the IPMI card is known to be good, and works in another system, it could indicate a problem with the +5V Standby as going through the motherboard, or coming from the power supply – in other words, a more serious problem.

CMOS Battery: We do show the status of the CMOS battery from the motherboard in the NumaRAID GUI. If the battery gets low (~6% of it's normal voltage), you will start to see symptoms of the battery failing, such as the date and time on the hardware clock are not correct, and bootup messages saying the battery is low or dead. It is very simple to replace and very low-cost. At the time of this writing, SuperMicro boards use CR-2032 3V batteries. Do NOT substitute other models, such as CR-2025.

SATA/SAS (On-board): We do use the on-board SAS/SATA controller(s) for our products. Some of the motherboards we use have up to 3 independent controllers – each different brands/models. Typically, on-board SATA is handled by the Intel ESB2 controller. If it fails, the array won't boot. Some other systems use Intel ICH-9R or ICH-10R RAID controllers. If the system is booting from this controller, and it fails, the system won't boot. Finally, some systems have a on-board LSI controller. If the boot drive is connected to this and it fails, the system won't boot. You can test the bootup by moving the boot drive to another system. SATA cables can also get damaged.

USB: Typically, USB ports are used for installation, but sometimes are also used for a keyboard or mouse. On some (rare) motherboards, the physical port used for the installation matters – this is because some motherboards have multiple USB chips. Also the built-in port enumerator might have a specific order for referencing the ports (Which is why from Linux, some ports appear to work better than others). Here's the problem with USB – it is delicate – just as delicate as the SATA connectors on the motherboard. It is really easy to snap off the plastic tab in the middle of the connector on the motherboard, so care must be taken when inserting or removing devices.

PS/2: While this is considered a legacy port, most motherboards still come with these connectors. They are very high-priority, in terms of interrupt, and are controlled (usually) by an Intel i8042 chip located somewhere on the motherboard. If this chip fails, both ports will go out.

CMOS/BIOS: If the BIOS dies, the motherboard is useless. However, if something is set incorrectly in the BIOS, it may prevent the array from operating properly. Motherboards with on-board RAID controllers may also have additional BIOSes for those – even a bootable Ethernet port might have it's own BIOS.

4.9 Drive Backplane problems

In general, there's two kinds of drive backplanes we use. One is a discreet backplane – the other is a SAS-switched backplane. Both types of backplanes have an SES2 enclosure management chip, which operates the LEDs and controls and monitors voltages, temperatures, and fans on the backplane itself. The way this chip connects to the host is different, however. On the switched backplanes, the chip is connected to the switch, whereas on non-switched backplanes, it connects to the host via an I2C interface. On the switched backplanes, the switch connects to the host via an I2C interface instead. How these backplanes are constructed varies: Typically, the discreet backplane has SAS connectors on the drives which go through the board (i.e. through hole), whereas on the switched backplane, the drive connectors are surface-mounted. Roughhousing the drives (i.e. not inserting them carefully) could damage the connectors. On the rear of the board, there are multilane connectors or discreet SATA connectors – these are also potentially very delicate. On the multilane connector, should the shield become bent, the cable may not seat properly, causing bad connections. Also, the I2C connection is especially delicate. Finally, there is power: Most of these boards have multiple power connections – this isn't done just to have a place to put the connectors – it's done for distributing the power across the ports – this enables hot-pluggability. If, for example, one power connection was used, then hot-plugging one drive might cause other drives to momentarily spin down then back up.

4.10 Boot device problems

The boot device does have some mortality – even if it is a SATADOM. Aside from an all-out failure, or power/cabling problems, something to watch out for is what happens when the boot drive is full. If the drive ever becomes 100% full, it will act as if it is read-only on bootup. This will cause a host of problems after bootup. The easy way out from this point is to clear the logs (NumaRAID and system).

4.11 Data Drive problems

Here is a list of errors we have experienced with data drives:

Drive won't spin up (Could be drive firmware or bad drive or power/interface problem).

Drive is clicking (Bad drive – indicates head alignment problem).

Drive spins up and down repeatedly (Indicates a failure of the drive tachometer on the spindle motor).

Drive responds but won't spin (Spindle motor failure).

SMART indicates a problem (Imminent failure of a drive component).

Slow drive (Could be start of head alignment problem).

Drive vibrating excessively (Spindle balance weight came off).

4.12 SAS HBA problems

The internal connections on the LSI or Supermicro SAS HBA can be damaged – especially the shielding on the multilane SAS connector. As mentioned before, if this shielding becomes bent, it may prevent the cable from locking in properly. But note how this card interfaces with everything: There are 8 lanes going from the PCIe slot on the motherboard into the SAS chip, and 8 lanes coming out of the chip going to the cables. There are a number of components on the board which can be damaged, which could cause a failure on a single SAS lane. There are (among others), 9 LEDs on the board – one LED (usually visible on the outside) is a heartbeat. This LED blinks to indicate that the processor on the board is functioning. If the BIOS on the

card gets screwed up, it won't blink. The 8 other LEDs show communication between the drives and the card. If one doesn't light, then chances are there is no communication on that port. Rechecking cables first is always the best thing. One other note: These cards typically use the LSI 1068e chip. This chip supports a maximum of about 192 devices. However the switched backplanes from SuperMicro don't have the same number of devices as the backplane itself. Backplanes up to 16 drives have a SAS chip which takes the space of 28 devices. The 24-drive backplane has a SAS chip which takes the space of 64 devices, so although the card supports 192 devices, using SuperMicro switched backplanes, it can't support more than (3) 24-drive backplanes or more than about (6) 16-drive backplanes. If you need more, instead of an LSI 3081e, use a variant called an LSI 3801e – It looks exactly the same, except instead of two 4-lane ports connecting to one channel, it has two 4-lane ports, each on a separate channel.

The internal discreet SATA connectors and especially a sideband connector –are especially delicate and prone to breakage. The actual card-edge connection portion of the multilane connector typically isn't a problem, what is, is the small metal spring button which secures the cable to the shield of the connector it is plugged into. This button can and will move or shift. When it's all the way back, towards the cable, the position will prevent it from locking into the shield – it must be all the way forward, and the two latches on it must lock to the shield in order to be sure that the card-edge connector on the cable is securely mated properly. If this latch becomes bent, it must be fixed – at all cost. If it can not be fixed, the cable has to be replaced. If the cable is used with a broken latch, then it's possible that not all of the drives connected to the cable will come up.

4.13 SAS Host connectivity issues

This is more of a tip than for troubleshooting. The cable is not very easy to damage. The main problem area is: "I can't get the cable out." At the front of the cable are two pairs of metal hooks which hook onto the socket. If you pull on the cable really hard, and pull on the release really hard, the cable might not come out – this is because you are trying too hard, and are actually pulling the hooks against the socket harder than the release is trying to release them. If this occurs, while holding the release on the cable, push the cable in (instead of out), and you will hear the latches release, then pull the cable out.

4.14 Fibre HBA problems

Note that this card is especially delicate – not so much in terms of ESD, but in regards to the physical components on the card. If the Fibre shields become damaged or distorted, it might not be possible to properly insert SFPs into them. Also on the back are a series of very tall surface-mount components (specifically some capacitors) – if these are broken off, specific ports won't work. These aside, single ports can fail, and multiple ports can fail. If all ports fail, try swapping the card, otherwise check the software, then the cables, then the SFPs. These small SFP is almost an entire computer in itself, with it's own PIC processor, RAM, signal noise filter, retimer, amplifier, laser diode, and optical detector. If any components in an SFP fail, it is not serviceable, and should be replaced. You can observe the output of the laser (carefully, but not too close). If there is no light, and the SFP is fully-inserted, either the device it is plugged into is not providing power, or the SFP is bad.


4.15 Fibre Host connectivity issues

Of all of the possible cables in Aurora LS's RAID system, by far, the most delicate are Fibre Channel cables. The amount of problems that is possible with these cables is somewhat astronomical compared with other cables. First a description of how they are constructed: There are two optical conduits in a standard LC cable, one carrying light to the array's SFP, and one bringing light back. The diameter of this conduit is much larger than the width of the laser beam projected into it, but the cable is designed to bounce the beam off the inner sides of the fibre conduit. At the ends of these conduits are a pair of lenses. These lenses are glued on carefully, by hand, and do two things: 1) Protect the ends of the fibre conduit itself, and 2) Focus the beam to a point going in or out. The lenses can occasionally get misaligned or move during the glue's curing process. Everything surrounding the lenses (just about) is plastic. Two mechanical problems which can occur are because of this plastic: It's possible that a cable may be misassembled by putting the wrong lens on the wrong conduit, flipping one end of the cable upside-down. The way to tell if this is occurring is to plug the cable into an operating Fibre channel device, and compare the other end to what it is plugging into – if the laser on what it is plugging into is coming from the same side as the laser coming from the cable, the cable is defective. The other mechanical problem – the plastic portion of the plugs can be broken easily, so care must be used when inserting or especially when removing the cables. Now the cable itself is made of fiberglass, which is essentially plastic. If you took a clear semi-thick piece of plastic and bent it, you would find that where it bends, it turns opaque (white), and you can't see through that part. It is similar with the fibre itself – you don't want to bend it if possible – I'd say you don't want to go around a bend with an equivalent diameter less than a 3 inch circle. If it is bent too far, although you won't be able to see it, the cable inside will turn opaque, preventing the beam from passing through properly. If this happens, the cable is useless. When the Fibre cables or cards are shipped, they have protective covers. The cover on the card is mainly to keep dust out (If dust gets in-between the emitter/detector and the lens, it might impair data transmission). However the covers on the cable are for a different reason – to protect the lenses from getting scratched. If the lenses on the cable become scratched, they will also impair the ability for the cable to carry the light from the laser.

4.16 Troubleshooting Aurora LS's Client Related Problems

Fibre Based Clients

Assuming there are no problems on the array, in order for a client to be able to see a LUN, there is a certain chain of items which must be present as in the following diagram:



Going clockwise from the upper left, you have to have a RAID in order to create a LUN. The user and Initiator are optional, however if there are any defined, the user you are trying to communicate with must also be set up as a user and given read only or read/write access to the LUN. Then the Target is optional, however if one is assigned to the LUN, then it must be on the same connection that the client is going to be connected to. The SFP on the array being used must be working – with no more than one connected to a switch if you are using a switch (unless you are doing some careful zoning on the switch). For troubleshooting, if you have a switch, you may want to remove it, otherwise either of the SFPs, cables, switch, or zoning could be the problem. If you are using a switch, and it is zoned, make sure the array and the client are in the same zone (I have had a tech support story once about a switch which was rented by a customer, and although they didn't zone the switch, the previous renter did, disabling the ports that were being used). Then on the client, the cable or SFP could be a problem, and the HBA could have a problem. There could be an OS problem (which is rare), or a problem with the driver. Here's the troubleshooting technique: If you look carefully at the chart, there is a straight chain, going from RAID to the Fibre Driver on the client. You should troubleshoot from one end of the chain to the other, otherwise it is confusing. Start by making sure there is a RAID, with a LUN on it. Next, look at Users, and see if the user is showing up at all. If not, skip to the other end of the chain, and start troubleshooting from that end. If the user is showing up under users, then it is almost certainly a problem with an Initiator or target setting – check to make sure either no targets exist, or that the target being used exists, and check to make sure the initiators exist, and that the user in question is assigned to that LUN, or that no initiators exist. If you had to troubleshoot going the other way, if the client is running OS/X, make sure the Fibre card/drivers are working by going into Apple System Profiler. If it is Linux, do an `lsmod` to find the Fibre driver. If it is Windows, go into the device manager, and make sure you can see the Fibre channel card under Storage devices, and that there is no yellow or red exclamation point next to it. If this is Linux, do an `lsccsi` to see if you can see the LUN. If it is Windows, go into Disk Management and see if you can see the LUN. If it is OS/X, go into Apple Disk Utility. At this point, if the array is all set correctly, and the client seems OK, you may have a hardware problem. Check the LEDs on the array and the client – they should

indicate a link at the speed of the client's adapter. If not, there might be a bad cable, SFP, or HBA.

4.17 Using IPMI to diagnose problems

Some NumaRAID arrays are equipped with an IPMI card (short for Intelligent Peripheral Management Interface). This card or on-board chip is literally a second computer, but is very small. It runs off the +5V standby which is used to power the on/off switch, and is capable of communicating through the motherboard even if the array is off. To access the IPMI services, make sure you have a connection to the IPMI Ethernet port. Change the TCP/IP settings on your client as follows:

IP address: 192.168.0.1

Subnet: 255.255.255.0

Gateway: 192.168.0.201

Open a network browser from the client, and type:


<http://192.168.0.201>[enter]

You should see a login screen which looks like the following:

A screenshot of a web-based login interface. At the top, the text "Authenticate with Login and Password!" is centered. Below this, there are two input fields: "Username" and "Password". Each field has a corresponding text label to its left. Below the password field is a blue "Login" button with white text.

The login screen will prompt for a user name and password. The defaults are ADMIN and ADMIN (must be capitalized).

The main IPMI window should appear as follows:



You can get into the IPMI card at any time.

If you left-click on Power Down, as in the image above, and the array is powered on, you won't be able to stop it – it will be off as if the power switch itself was pressed. Also, if you left-click on the Reset button, the array will reboot as if you actually hit the reset button on the front.

You can turn on the array via the power on button. Once the array is on and starting to boot, you can click on the small window in the middle and bring up the console as if you were actually looking at it on the monitor. This is the primary area where you may have to troubleshoot the array. You can even view or control the BIOS from here. On the left, each of these items is a menu which expands downward if you click on them. If you are troubleshooting power problems, the main item that you want is System Health. Once you've left-clicked on this, you can click on "Monitor Sensors."

Monitor Sensors will bring up the following window:

Monitoring Sensors

Sensor Type	Sensor Name	Sensor Status	Sensor Reading
Temperature	CPU Temp 1	OK	33 degrees C
Temperature	CPU Temp 2	OK	32 degrees C
Temperature	CPU Temp 3	No reading	
Temperature	CPU Temp 4	No reading	
Temperature	Sys Temp	OK	39 degrees C
Voltage	CPU1 Vcore	OK	1.184 (+/- 0.004) Volts
Voltage	CPU2 Vcore	OK	1.192 (+/- 0.004) Volts
Voltage	3.3V	OK	3.280 Volts
Voltage	5V	OK	4.920 (+/- 0.012) Volts
Voltage	12V	OK	11.904 (+/- 0.048) Volts
Voltage	-12V	OK	-12.100 (+/- -0.050) Volts
Voltage	1.5V	OK	1.504 (+/- 0.008) Volts
Voltage	5VSB	OK	4.920 (+/- 0.012) Volts
Voltage	VBAT	OK	3.392 (+/- 0.008) Volts
Fan	Fan1	OK	8400 RPM
Fan	Fan2	OK	8700 RPM
Fan	Fan3	OK	8500 RPM
Fan	Fan4	Below lower non-recoverable threshold	0 RPM
Fan	Fan5	Below lower non-recoverable threshold	0 RPM
Fan	Fan6	Below lower non-recoverable threshold	0 RPM
Fan	Fan7/CPU1	Below lower non-recoverable threshold	0 RPM
Fan	Fan8/CPU2	OK	8700 RPM
Physical Security	Intrusion	OK	
Power Supply	Power Supply	OK	
Module / Board	CPU0 Internal E	OK	
Module / Board	CPU1 Internal E	OK	
Module / Board	CPU Overheat	OK	
Module / Board	Thermal Trip0	OK	
Module / Board	Thermal Trip1	OK	
System Firmware Progress	BIOS	OK	

Refresh

You are interested in the items which are red on the left. In the example above, four of the fans are in red – this was normal for this particular model, where there are no fans connected to connectors 4 through 7. On this system, there were (5) fans – one each on connectors 1 through 3 and connectors 7 and 8.

Back on the main window, which shows the remote desktop, you should be able to diagnose a problem if the computer isn't booting. It should indicate an error on the screen. Common errors might be due to a faulty boot drive, a disk in the DVD-ROM drive, or even a USB device.

If you are troubleshooting a network problem, when the computer boots to the first Red Hat logo screen with choices, immediately hit an arrow key to stop it from counting, then select Safe Boot. Watch the screen carefully, it will eventually reach a login prompt, login as root

(default password is rdserdse). Once you are in, you can type `ifconfig` and look at the network settings, and change them with the `system-config-network-tui` command.



CAUTION: It is very important that when you are not using the IPMI interface via the web browser, that you click the logout icon on the upper right before you exit your browser.


Section 5

Application / Technical / Customer Notes


5.0 Application / Technical / Customer Notes

5.1 Windows Infiniband Performance Tuning


In Windows, you can improve performance. However to do so, you will need to edit the registry using a program called “regedit.” Left-click on the Windows logo (or Start button):




Left-click on All Programs (or Program Files):




Left-click on Accessories:



Left-click on Command Prompt:




This will open a command prompt window, which looks as follows:




From the command prompt, type:

regedit[enter]


Once you run regedit, drag the scrollbar on the left area all the way to the top if it is not already:




Left-click on Computer in the upper left corner of the left window (You are going to do a search, this starts the search at the very beginning):




Type: **[Ctrl-F]**. This will open the search window. The text box which is prompting what to search for is already selected, type: **ModeFlags** (Already typed in the picture):




At the bottom of the search window, make sure at the bottom, that Look at Values and Match whole string only are selected:




Left-click on the Find Next button:



The computer will search for the text specified. When it finds something, the “Searching” window will disappear, and you will see the entry on the right as follows:




Press **[enter]**. This will cause a small pop-up window to appear:



Press: **[2][enter][f3]**. This will change the value, close the window, and continue searching.

You will know the search is complete when the following pop-up window appears:




When it is finished, left-click on the OK button, and close the Registry Editor, then reboot/restart the client. This yields a speed increase up to 30%.

5.2 Additional Administration Functions

Webmin has a number of additional functions which can provide additional functionality to the Aurora. The functions listed here are the ones specific to NumaRAID – other functions can be dangerous, and are not discussed.

System Information



Login: root

- Webmin
- System
- Servers
- Others
- Networking
- Hardware
- Cluster
- Un-used Modules

Search:

[View Module's Logs](#)

[System Information](#)

[Refresh Modules](#)

[Logout](#)

System hostname aurora

Operating system Redhat Enterprise Linux 5

Webmin version 1.441

Time on system Thu Sep 24 10:34:47 2009

Kernel and CPU Linux 2.6.18-92.el5 on x86_64

CPU load averages 0.16 (1 min) 0.25 (5 mins) 0.17 (15 mins)

Real memory 7.67 GB total, 3.06 GB used

Virtual memory 767.99 MB total, 0 bytes used


Local disk space 6.51 GB total, 2.51 GB used

The main Webmin System Information screen provides some information. It is either the first screen you see after logging into Webmin, or in the webmin menu on the left, you can left-click on [System Information](#) located near the bottom of the menu. Items which would be of interest are the System hostname, which shows the name of the array, Time on system indicates the current date/time on the array. Real memory shows how much physical memory is available to the operating system, and how much is free. Local disk space shows the total capacity of the boot device, and how much is used.

To return to NumaRAID functions, expand the Hardware group on the left, if it is not already, then left-click on NumaRAID GUI under this group.

IP Address Firewall

It is possible to set webmin to deny or allow specific IP addresses access to the array. To do this, expand the Webmin group on the left, then click on Webmin Configuration under this group. A series of icons will appear on the right, as follows:



Left-click on the icon which reads IP Access Control. This brings up the following screen:



Notice the bubbles at the top of the table. If you check the left bubble (the default) – Allow from all addresses, All IP addresses will be able to access this array. The other two bubbles are used in conjunction with the text box below. You enter IP addresses into the text box. If you then check the bubble which reads Only allow from listed addresses, then only IP addresses listed in the text box will be able to access this array. If you check the right bubble Deny from listed addresses, then any IP address except the ones listed will be able to access this array. Once you have the screen set the way you would like, left-click on the Save button at the bottom. To exit without saving, left-click on the Return to Webmin configuration link at the bottom of the screen.

To return to the NumaRAID GUI, expand the Hardware category on the left, if it is not already, and left-click on NumaRAID GUI.

Default to NumaRAID GUI after Login

You can set Webmin, so that the NumaRAID GUI is always the first thing which appears after login by doing the following:

Expand the Webmin Group if it is not already, so that you can see the items under it.

Left-click on the Webmin Configuration item below the Webmin Group.

On the right, left-click the icon which reads “Index Page Options.”

Near the bottom of the table is a line which reads “After login, always go to module” – next to this is a drop-down. Left-click on the drop-down, and select NumaRAID GUI.

Left-click on the Save button at the bottom of the screen.

To return to the NumaRAID GUI, expand the Hardware category on the left, if it is not already, and left-click on [NumaRAID GUI](#).

Make the NumaRAID GUI a Little Faster

You can make the NumaRAID GUI a little bit faster by forcing Webmin to cache it's libraries. To do this, do the following:

Expand the Webmin Group if it is not already, so that you can see the items under it.

Left-click on the Webmin Configuration item below the Webmin Group.

On the right, left-click the icon which reads “Advanced Options.”

The fourth item down reads “Pre-load Webmin functions library.” Select the bubble next to Yes.

Left-click on the Save button at the bottom of the screen.

To return to the NumaRAID GUI, expand the Hardware category on the left, if it is not already, and left-click on [NumaRAID GUI](#).

Find the IP Addresses of Other NumaRAID(s) on the Network

If you are logged into one NumaRAID array, and want to find out the addresses of other NumaRAID arrays on the same network, you can do the following:

Expand the Webmin group on the left, if it is not already.

Under the Webmin group, left-click on [Webmin Servers Index](#).

At the top, left-click on the button which reads [Broadcast for servers](#).

To return to the NumaRAID GUI, expand the Hardware category on the left, if it is not already, and left-click on [NumaRAID GUI](#).

Adding/Deleting/Changing Webmin Users

You can create other users/logins for Webmin, without having to create Linux users. This is done as follows:

Expand the Webmin group on the left, if it is not already.

Under the Webmin group, left-click on [Webmin Users](#).

At the top, you will see a table of users. You can either left-click on the link above or below the table, which reads [Create a new Webmin User](#), if you would like to create a user. To delete a user, you can left-click to turn on the checkbox next to the user name, then left-click on the [Delete Selected](#) button at the bottom. You can edit information for a user by just left-clicking on their [user name](#).

If you are creating a user, the screen will change. Enter the Username and password at the top, then scroll down to the bottom and left-click on the [Create](#) button.

To return to the NumaRAID GUI, expand the Hardware category on the left, if it is not already, and left-click on [NumaRAID GUI](#).

Changing Passwords

If you want to only change the webmin password for a user, follow the process above. You can, however, change the Linux password for a user, by doing the following:



CAUTION: If you change the root password, then later forget this password, you will have a serious problem, as you may not be able to log back in to make changes.

Expand the System group on the left, if it is not already.

Left-click on the option within this group, called Change Passwords.

Left-click on the user name whos password you would like to change.

Type the new password (twice).

Make sure the Force user to change password at next login? option is NOT checked, and make sure the Change password in other modules? option IS checked, then left-click on the Change button.

To return to the NumaRAID GUI, expand the Hardware category on the left, if it is not already, and left-click on NumaRAID GUI.

Run a CLI command from Webmin

It is possible to run CLI commands from Webmin. To do so, do the following:

Expand the “Others” group on the left if it is not already.

Left-click on Command Shell below the Others group.

Type the command you would like, and press **[enter]**.

To return to the NumaRAID GUI, expand the Hardware category on the left, if it is not already, and left-click on NumaRAID GUI.

Change the Network Host Name

To change the network host name, do the following:

Expand the Networking group on the left if it is not already.

Left-click on Network Configuration.

The screen on the right will change, left-click on Hostname and DNS Client.

At the top, change the hostname.

At the bottom, left-click on the Save button.

To return to the NumaRAID GUI, expand the Hardware category on the left, if it is not already, and left-click on NumaRAID GUI.

See and Control SMART for the Boot Device

You can see the status of SMART for the boot device, as well as run SMART diagnostic tests on it, by doing the following:

Expand the Hardware group on the left if it is not already.

Left-click on SMART Drive Status.

At the top of the screen, make sure the boot drive is selected.

Left-click on the Show button.

Once you have seen the data, you will get option buttons at the bottom to run a Short Self-Test, Extended Self-Test, or a Data Collection test. These tests are not destructive, and will not result in any loss of data. Note that the extended test can take a long time to run, during which time, the array will be inaccessible.

To return to the NumaRAID GUI, expand the Hardware category on the left, if it is not already, and left-click on [NumaRAID GUI](#).

Setting System Time or Timezone

Over time, you may find that the time/date on the array is not accurate, and may need to be occasionally adjusted. Also, the time zone might not match your location. There are two clocks in the system. One clock is the hardware clock, the other is a system (software) clock. The system clock reads the hardware clock when it is first booted, then after that the system clock is mathematically calculated as an offset using the system timer. The accuracy of this timer can drift, and the system clock may not match the hardware clock over time. The hardware clock can also drift. To get to the time screen, do the following:

Expand the Hardware group, if it is not already expanded.

Left-click on System Time. On the right, the following screen will appear:

Help...
Module Config

Search Docs...

Set time Change timezone Time server sync

This form is for changing the system's current time, which is used by all running processes. On operating systems that have a separate hardware clock, it can be used to set that too.

System Time

Day	Date	Month	Year	Hour, minute and second
Tuesday	29	September	2009	15 : 10 : 26

Apply Set system time to hardware time

Hardware Time

Day	Date	Month	Year	Hour, minute and second
Tuesday	29	September	2009	15 : 10 : 26

Save Set hardware time to system time

If you wish to change the timezone, left-click on the [Change timezone](#) tab at the top of the screen, then change the timezone, and left-click on the [Save](#) button at the bottom.

On this screen, you can set the system time, hardware time, or both. Set the time and/or date using the drop-downs. But here's the gist on the buttons. Under system time is an Apply button. This is used to set the (software) system time. It isn't a save button, because the software/system time isn't saved anywhere – it is just an offset running from RAM. The Set system time to hardware time button will set the system/software time to the current time read from the hardware clock. In the lower table, is a Save button. This is used to save the current hardware time. This is set in non-volatile memory inside the array. The Set hardware time to system time button sets the hardware time to the current system/software time.

To return to the NumaRAID GUI, expand the Hardware category on the left, if it is not already, and left-click on [NumaRAID GUI](#).

Logging Out

Although you do not have to log out of the array, it is better if you do, as the logging in/out are logged by Webmin. To logout, simply left-click on [Logout](#) at the bottom of the left menu.

5.3 Fibre Channel Switch Zoning

With Fibre Channel, complexities are shifted from the client to the switch (if you are using a switch). On the client, no more software is required other than the driver for the Fibre Channel HBA, and the operating system itself. This section will focus more on switch zoning concepts.

While it might seem that you can just take a Fibre channel switch out of the box, just plug it in, and use it, this is not always the case. By default, a Fibre channel switch usually will act as a hub, but a very powerful one. Any data coming in from any of the ports (by default) will be sent to all of the other ports simultaneously. But this is not a good thing: Arrays will try to send data to arrays, clients will try to send data to clients, and arrays and clients which were not intended to communicate with each other will communicate. From an array management standpoint, while this might not be a problem, it creates a lot of unnecessary traffic on the switch and everything connected to it, which can have an adverse effect on data rates.

Earlier switches, such as 1Gbit switches and some 2Gbit switches used a technique called “provisioning” to govern the connections, however it wasn’t very efficient from a management point of view (or lack thereof). It worked like this: The clients are called initiators, and arrays are called targets. You would flash the firmware on the switch, such that a certain number of ports are allocated for initiators, and a certain number are allocated for targets. This would prevent the problems with clients communicating with clients, and arrays communicating with arrays, but it still didn’t fix the problem with clients communicating with unintended arrays and vice-versa.

Newer switches are called “fabric switches,” and use what is called “zoning” instead of provisioning. The term fabric is referring to a meshed grid which is formed by initiators and targets, with the initiators and targets on the “fringe” of the grid. This is more advanced, and solves all problems, however there is a lot of thinking which is involved, and the software can be quite involved. The key to zoning is being able to mentally visualize the setup.

At it’s simplest, a zone is a fabric “bag” which contains ports. You can usually zone the switches in such a way, that you can have any number of zones, with any number of ports in each one, and they can overlap. The zone does not differentiate between an initiator or target – they are just connection points. So, using the bag as an example, and a mechanical nut for a target, and bolt for an initiator. You can have a bag of nuts, a bag of bolts, or a bag of nuts and bolts.

In more complex setups, you want to avoid the pitfall of creating “mini-provisioning” problems – to make this easy, you don’t want to have more than one nut or bolt in the same bag unless there’s no alternative – or in real life, no more than 2 ports in one zone.

In general, you only want two-way communication – not 3 or more way communication. Overlapping zones are OK, as long as they are thought out. As switch connections increase, it may be necessary to have more than one switch – this is called cascading switches. The problem is that cascading the switches kind of goes back to provisioning, in that you can/should only have one cascade port in a zone – without any others (unless they are in other zones). A certain 4Gbit switch I know of has (24) 4Gbit ports, but has (4) 10Gbit cascade ports. You obviously won’t get the throughput of 48Gbits coming to/from the 4Gbit ports going at the same speed going through the 10Gbit ports, so careful planning has to be done when scaling up with switches.

5.4 Infiniband Switch Configurations

Quick note about Infiniband switches: Infiniband switches are not the same as Fibre Channel switches, because of how the subnet is run. The subnet is run by clients, so switch management or zoning isn't necessary in most cases – the switches are just for connecting single or multiple arrays to single or multiple clients